

An asymptotic optimality result for the multiclass queue with finite buffers in heavy traffic*

Rami Atar

Mark Shifrin

Department of Electrical Engineering
Technion–Israel Institute of Technology
Haifa 32000, Israel

August 21, 2013; revised May 5, 2014

Abstract

For a multiclass G/G/1 queue with finite buffers, admission and scheduling control, and holding and rejection costs, we construct a policy that is asymptotically optimal in the heavy traffic limit. The policy is specified in terms of a single parameter which constitutes the free boundary point from the Harrison-Taksar free boundary problem, but otherwise depends “explicitly” on the problem data. The *cμ priority rule* is also used by the policy, but in a way that is novel, and, in particular, different than that used in problems with infinite buffers. We also address an analogous problem where buffer constraints are replaced by throughput time constraints.

AMS subject classifications: 60F17, 60J60, 60K25, 93E20

Keywords: Multiclass G/G/1 queue; Brownian control problems; Bellman equation; The Harrison-Taksar free boundary problem; State dependent priorities

1 Introduction

In this work we consider the problem of finding asymptotically optimal (AO) controls for the multiclass G/G/1 queue with finite buffers, in heavy traffic. Upon arrival of a class- i customer into queue i (with $i \in \{1, \dots, I\}$ and where I denotes the number of classes), a decision maker may either accept or reject the job. In addition, the decision maker controls the fraction of effort devoted by the server to the customer at the head of queue i , for each i . We refer to the two elements of the control as *admission control* and *scheduling*, respectively. The problem considered is to minimize a combination of holding and rejection costs. The term ‘heavy traffic’ refers to assuming a critical

*Research supported in part by the ISF (Grant 1349/08), the US-Israel BSF (Grant 2008466), and the Technion fund for promotion of research

load condition and observing the model at diffusion scale. Our interest in this problem stems from recent developments in the application area of cloud computing. In a hybrid cloud where a private cloud (namely, a local server) has a given capacity and memory limits, tasks that cannot be queued in real time are rejected from the local system and sent to a public cloud, where a fixed charge per usage applies. For further details on modeling toward these applications, see [38]. For a more general modeling framework of data centers, see [12]. The analysis of the model leads in the scaling limit to a control problem associated with Brownian motion (BM), often referred to in this context as a *Brownian control problem* (BCP). Our main result is the convergence of the *queueing control problem* (QCP) value function to that of the BCP, and the construction of a particular AO admission/scheduling policy. The policy is specified in terms of a free boundary point that is used in solving the BCP, but otherwise it depends explicitly on the problem data.

A line of research starting from Harrison [25] and continuing with Harrison and van Mieghem [29], Harrison [26], [27] and Harrison and Williams [30] has treated BCP associated with a broad family of models called *stochastic processing networks*. These problems, aimed at describing the heavy traffic limits of QCP, were shown to be equivalent to *reduced BCP* (RBCP), in which *workload* plays the role of a state process. RBCP simplify BCP in two ways: Their state lies in lower dimension, and their form, specifically, that of a singularly controlled diffusion, makes control theoretic tools applicable. Addressing these models at the same level of generality, Atar and Budhiraja [6] and Atar, Budhiraja and Williams [7] use such control theoretic tools to characterize the BCP (equivalently, RBCP) value functions as solutions to *Hamilton-Jacobi-Bellman* (HJB) equations, and Budhiraja and Ghosh [15] and [16] prove convergence of QCP value function to BCP value functions. Many other works address these models in situations where the BCP are explicitly solvable, see e.g. Ata and Kumar [2] and references therein.

As far as BCP are concerned, the model studied here is a special case of the models considered in some of the aforementioned papers. In particular, BCP and RBCP play here important roles, where the reduction from an I -dimensional BCP to a one-dimensional RBCP is a special case of [30]. Moreover, the HJB equation, that in the present setting is an ordinary differential equation and will be referred to merely as a *Bellman* equation, is a special case of the partial differential equations treated in [7]. In addition, our specific one-dimensional RBCP, its relation to the Bellman equation, and its solution go back to Harrison and Taksar [28], where a singular control problem for a BM is solved. The solution is given by a reflected BM (RBM), with supporting interval determined by a free boundary problem associated with the Bellman equation. This type of free boundary problem first appeared in [28], and we therefore refer to it as the *Harrison-Taksar free boundary problem*. In our case, the interval is always of the form $[0, \mathbf{x}^*]$, and we call \mathbf{x}^* the free boundary point.

On the other hand, the works [15], [16] and [2], despite their vast generality, do not cover the present model as they do not treat admission control and rejection penalties. Thus, while the BCP is well understood, convergence and AO issues have not been addressed before. Addressing these issues is the main contribution of this paper. This is done by proving that the BCP value function constitutes a lower bound on the limit inferior of QCP costs under any sequence of policies (Theorem 3.1), and then constructing a specific policy that asymptotically achieves this lower bound (Theorem 4.1). This AO policy depends explicitly on the system parameters, except that it also depends on the quantity \mathbf{x}^* . Moreover, it uses the well-known *cμ rule* in a novel way, as we now explain.

The structure of the policy alluded to above is simple enough to describe without introducing much notation. The notation needed is as follows. For class- i customers, denote holding cost per unit time by h_i , rejection penalty per customer by r_i , and reciprocal mean service time by μ_i . The policy is defined in terms of three elements: The index $h_i\mu_i$, the index $r_i\mu_i$, and the free boundary point \mathbf{x}^* . The first index is used for scheduling. It is precisely the index used for the $c\mu$ priority rule (a terminology used when c , rather than h , denotes holding cost per unit time), where classes are prioritized in the order of $h_i\mu_i$, the highest priority given to the class i with greatest $h_i\mu_i$. As observed first by Smith [39] and Cox and Smith [18], the $c\mu$ priority rule is exactly optimal for holding costs. Many extensions to this result have been shown (see e.g., [17], [40] and discussions therein). Our scheduling policy uses the same index to assign priorities, but in a state-dependent fashion, as follows. At any given time, the lowest priority is assigned to the class i having lowest index $h_i\mu_i$ among classes for which the buffers are not nearly full. We give precise meaning to the term ‘nearly full’.

Let us contrast this with the case of infinite buffers and no rejection. For this model, an AO policy applying dynamic priorities, in the form of an extended version of the $c\mu$ rule, was developed by van Mieghem [40] to address nonlinear delay costs. When costs are linear, as they are in the present paper, it is the *fixed* priority rule according to $h_i\mu_i$ that is AO. Suppose now that I is the class that has lowest $h_i\mu_i$ value, so that class I is assigned lowest priority by this rule. Then, as is well-known since Whitt [42], the multiclass G/G/1 queue behaves in such a way that all classes $i < I$ exhibit vanishing queue length in the heavy traffic limit. Consequently, it is not only the aforementioned assignment rule that is AO. Any priority policy assigning lowest priority to the class I performs equally well, and is therefore AO for such a QCP. In other words, the only aspect of the index policy which is important for AO in the problem with infinite buffers and no rejections, is the class assigned the lowest priority. Thus there is a major difference between the way in which the index is used in the infinite buffer setting and in this paper. In the latter case, the full information on the ordering of classes is important.

The admission control is based on the other index, $r_i\mu_i$, and the free boundary point \mathbf{x}^* . The significance of this index for admission control in heavy traffic was first noticed by Plambeck, Kumar and Harrison [34] (see below). Our policy acts as follows. When the diffusion-scaled workload level exceeds the level \mathbf{x}^* , all arrivals of one particular class are rejected. This is the class i having the least $r_i\mu_i$ value. When the workload level is below \mathbf{x}^* , all arrivals are admitted, except rejections that must take place so as to keep the buffer size constraint valid (namely arrivals that occur at a time when the corresponding buffer is full). We call these *forced rejections*. A property of the policy that is important for AO is that it maintains, with high probability, a low number of forced rejections. As a result, nearly all rejections occur when the workload exceeds \mathbf{x}^* , and only from one class. It is to this end that the scheduling policy prioritizes classes with nearly full buffer.

The aforementioned paper [34] studies the problem of minimizing rejection penalties, subject to throughput time constraints, for the multiclass G/G/1 queue in heavy traffic (see also Ata [1] for a closely related formulation). Each class has a deterministic constraint on the throughput time, and arrivals that are admitted into the system are assured that, with high probability, their throughput time constraint will be kept. This property of the policy is referred to as *asymptotic compliance*. The policy of [34] admits all arrivals except those from the class having lowest $r_i\mu_i$ value, and only when the workload exceeds a threshold value. Thus our admission policy resembles that of [34], except that our threshold level is characterized by the free boundary problem, whereas it is explicit

in [34] (their scheduling policy is different than ours).

But the relation of our work to [34] is deeper than similarities in the admission policies. Reiman's snapshot principle [36], and the pathwise Little's law, state that, under suitable assumptions, a deterministic relation holds in the heavy traffic limit between throughput time and queue length processes. Accordingly, buffer constraints on queue length should be asymptotically equivalent to throughput time constraints. We follow this rationale in the last section of this paper, where we formulate a QCP that parallels the QCP addressed in the main body of the paper, where finite buffer constraints are replaced by throughput time constraints. This may be regarded an extended version of the QCP of [34] that accommodates holding costs. We do not succeed in fully solving this problem here; our purpose in this part of the work is mainly to pose the problem and to discuss similarities with the main body of this paper, leaving the main question open. We begin by proving a pathwise Little's law in the form of a conditional result (Proposition 5.1). There is no guarantee that queue length and throughput times satisfy Little's law under an arbitrary sequence of controls. We show that C -tightness of the processes involved suffices. Using this result we can show that the policy we develop for the finite buffer problem satisfies the throughput time constraints and that its limit performance is dominated by the BCP value (Theorem 5.1). In order to deduce that it is AO, a lower bound in the same form is also needed. However, due to the lack of validity of Little's law for general sequences of policies, we can only show AO in a restricted class of policies (Proposition 5.2). The broader problem, and hence the question of AO remain open (see Conjecture 5.1).

Under the AO policy, the I -dimensional queue length process converges to the process solving the BCP. This convergence is a form of a *state space collapse* (SSC), a term referring to a behavior where queue length process limits are dictated by workload process limits. SSC is an important ingredient in the analysis of queueing network models in heavy traffic. It has been considered in many works, and in particular in a general setting by Bramson [14] and Williams [43]. The form of the SSC obtained in this paper involves spatial inhomogeneity due to the dynamic priorities, and is not covered by [14], [43], or, to the best of our knowledge, any other work on SSC. A part of the proof of Theorem 4.1 is aimed at showing a SSC result.

For a different formulation of a QCP with finite buffers and rejection costs, see Ghosh and Weerasinghe [24]. For a formulation other than [34] that combines asymptotic compliance and asymptotic optimality see Plambeck [35]. See Ward and Kumar [41], Rubino and Ata [37] and Ata and Olsen [3] for other treatments of AO in heavy traffic via a Bellman equation with free boundary, and Dai and Dai [19] for results on heavy traffic for systems with finite buffers without optimal control aspects. Finally, see Ghamami and Ward [23] for asymptotic optimality results based on a Bellman equation for the BCP, for a model with customer abandonment rather than rejection.

We will use the following notation. Given $k \in \mathbb{N}$, $\{e^{(i)}, i = 1, \dots, k\}$ denote the standard basis in \mathbb{R}^k . For $x \in \mathbb{R}$, $x^+ = \max(x, 0)$. For $a, b \in \mathbb{R}^k$, $a = (a_i)_{i=1, \dots, k}$, $b = (b_i)_{i=1, \dots, k}$, we denote $\|a\| = \sum_{i=1}^k |a_i|$ and $a \cdot b = \sum_{i=1}^k a_i b_i$. For $y : \mathbb{R}_+ \rightarrow \mathbb{R}^k$ and $T > 0$, $\|y\|_T = \sup_{t \in [0, T]} \|y(t)\|$. The modulus of continuity of y is given by

$$\bar{w}_T(y, \theta) = \sup\{\|y(s) - y(t)\| : s, t \in [0, T], |s - t| \leq \theta\}, \quad \theta, T > 0. \quad (1)$$

For Polish space E , denote by $\mathcal{D}_E[0, T]$ the space of RCLL maps from $[0, T]$ to E , equipped with

the usual Skorohod topology. A sequence of stochastic processes with sample paths in this space is said to be *C-tight* if it is tight and every subsequential limit has continuous sample paths w.p.1. Convergence in distribution of a sequence of random variables $\{X_n\}$ to X is denoted by $X_n \Rightarrow X$. For $a, b \in \mathbb{R}$, $a < b$, $C^2[a, b]$ denotes the set of functions from $[a, b]$ to \mathbb{R} that are twice continuously differentiable on (a, b) , for which derivatives of order ≤ 2 have continuous extensions to $[a, b]$.

The rest of this paper is organized as follows. In the next section we introduce the queueing model and QCP. Then we formulate the BCP and the RBCP, and state their solution via the Harrison-Taksar free boundary problem. We then discuss the interpretation of the solution. Section 3 shows that the BCP value function is a lower bound on the limit inferior of the sequence of value functions for the QCP. Section 4 constructs a policy for the QCP and proves that it is AO. Section 5 proves pathwise Little's law and relates the main body of the paper to the throughput time constraints formulation of [34].

2 Queueing and diffusion models

2.1 The multiclass G/G/1 model

A sequence of systems is considered, indexed by $n \in \mathbb{N}$. Quantities that depend on n have n as superscript in their notation. The system has a single server and $I \geq 1$ buffers, where each buffer is dedicated to a class of customers. The capacity of each of the buffers is limited, where the precise formulation of capacity is presented later. Customers that arrive at the system may either be accepted or rejected. Those that are accepted are queued in the corresponding buffers. Within each class, service is provided in the order of arrival, where the server only serves the customer at the head of each line. Processor sharing is allowed, in the sense that the server is capable of serving up to I customers (of distinct classes) simultaneously. An *allocation vector*, representing the fraction of effort dedicated to each of the classes, is any member of

$$\mathcal{B} := \left\{ \beta \in \mathbb{R}_+^I : \sum_{i \in \mathcal{I}} \beta_i \leq 1 \right\},$$

where, throughout, $\mathcal{I} = \{1, 2, \dots, I\}$.

A probability space $(\Omega, \mathcal{F}, \mathbb{P})$ is given, on which all random variables and stochastic processes involved in describing the model will be defined. Expectation w.r.t. \mathbb{P} is denoted by \mathbb{E} . Arrivals occur according to independent renewal processes. Let parameters $\lambda_i^n > 0$, $i \in \mathcal{I}$, $n \in \mathbb{N}$ be given, representing the *reciprocal mean inter-arrival times* of class- i customers in the n -th system. Let $\{IA_i(l) : l \in \mathbb{N}\}_{i \in \mathcal{I}}$ be independent sequences of strictly positive i.i.d. random variables with mean $\mathbb{E}[IA_i(1)] = 1$, $i \in \mathcal{I}$ and squared coefficient of variation $\text{Var}(IA_i(1))/\mathbb{E}[IA_i(1)]^2 = C_{IA_i}^2 \in (0, \infty)$. With $\sum_1^0 = 0$, the number of arrivals of class- i customers up to time t , for the n -th system, is given by

$$A_i^n(t) = A_i(\lambda_i^n t), \quad \text{where} \quad A_i(t) = \sup \left\{ l \geq 0 : \sum_{k=1}^l IA_i(k) \leq t \right\}, \quad t \geq 0. \quad (2)$$

The parameters λ_i^n satisfy

$$\lambda_i^n = n\lambda_i + \sqrt{n}\hat{\lambda}_i + o(\sqrt{n}), \quad (3)$$

where $\lambda_i > 0$ and $\hat{\lambda}_i \in \mathbb{R}$ are fixed.

Similarly, let parameters $\mu_i^n > 0$, $i \in \mathcal{I}$, $n \in \mathbb{N}$ be given, representing *reciprocal mean service times* for service to class i in the n -th system. Let independent sequences $\{ST_i(l) : l \in \mathbb{N}\}_{i \in \mathcal{I}}$ of strictly positive i.i.d. random variables (independent of the sequences $\{IA_i\}$) be given, with mean $\mathbb{E}[ST_i(1)] = 1$ and squared coefficient of variation $\text{Var}(ST_i(1))/\mathbb{E}[ST_i(1)]^2 = C_{ST_i}^2 \in (0, \infty)$. The time required to complete the l -th service to a class- i customer in the n -th system is given by $ST_i(l)/\lambda_i^n$ units of time dedicated by the server to this class. This can otherwise be stated in terms of the *potential service time* processes, given by

$$S_i^n(t) = S_i(\mu_i^n t), \quad \text{where} \quad S_i(t) = \sup \left\{ l \geq 0 : \sum_{k=1}^l ST_i(k) \leq t \right\}, \quad t \geq 0. \quad (4)$$

$S_i^n(t)$ is the number of class- i jobs completed by the time when the server has dedicated t units of time to work on jobs of this class. It is assumed that μ_i^n satisfy

$$\mu_i^n = n\mu_i + \sqrt{n}\hat{\mu}_i + o(\sqrt{n}), \quad (5)$$

where $\mu_i > 0$ and $\hat{\mu}_i \in \mathbb{R}$ are fixed. The first order quantities λ_i and μ_i are assumed to satisfy the *critical load condition*

$$\sum_{i \in \mathcal{I}} \rho_i = 1, \quad \text{where} \quad \rho_i = \frac{\lambda_i}{\mu_i}, \quad i \in \mathcal{I}. \quad (6)$$

The number of class- i rejections until time t in the n -th system is denoted by $Z_i^n(t)$. Since rejections occur only at times of arrival, we have

$$Z_i^n(t) = \int_{[0,t]} z_s^{n,i} dA_i^n(s) \quad (7)$$

for some process $z^{n,i}$.

The number of class- i customers present in the n -th system at time t is denoted by $X_i^n(t)$. For simplicity, the initial number of customers, $X_i^n(0)$ is deterministic, and it is assumed that no partial service has been provided to any of the jobs present in the system at time zero. We will call $X^n = (X_i^n)_{i \in \mathcal{I}}$ the queue length process. Let $B^n = (B_i^n)_{i \in \mathcal{I}}$ be a process taking values in the set \mathcal{B} , representing the fraction of effort devoted by the server to the various customer classes. Then

$$T_i^n(t) = \int_0^t B_i^n(s) ds \quad (8)$$

gives the time devoted to class- i customers up to time t . The number of service completions of class- i jobs during the time interval $[0, t]$ can thus be expressed in terms of the potential service process and the process T_i^n as

$$D_i^n(t) = S_i^n(T_i^n(t)). \quad (9)$$

We thus have

$$X_i^n(t) = X_i^n(0) + A_i^n(t) - D_i^n(t) - Z_i^n(t) = X_i^n(0) + A_i^n(t) - S_i^n(T_i^n(t)) - Z_i^n(t), \quad t \geq 0. \quad (10)$$

Use the notation A^n for $(A_i^n)_{i \in \mathcal{I}}$ and similarly for the processes S^n, Z^n, D^n, X^n, T^n . It is assumed that B^n has RCLL sample paths. By construction, the arrival and potential service processes also have RCLL paths, and accordingly, so do D^n, Z^n and X^n .

We define a rescaled version of the processes at diffusion scale as

$$\begin{aligned}\hat{A}_i^n(t) &= \frac{A_i^n(t) - \lambda_i^n t}{\sqrt{n}}, & \hat{S}_i^n(t) &= \frac{S_i^n(t) - \mu_i^n t}{\sqrt{n}}, & i \in \mathcal{I}, \\ \hat{Z}^n(t) &= \frac{Z^n(t)}{\sqrt{n}}, & \hat{X}^n(t) &= \frac{X^n(t)}{\sqrt{n}}.\end{aligned}\tag{11}$$

We now come to the buffer structure. A bounded closed convex set with nonempty interior $\mathcal{X} \subset \mathbb{R}_+^I$ is given, satisfying $0 \in \mathcal{X}$. It is assumed that, for every n , the rescaled initial condition $\hat{X}^n(0)$ lies in \mathcal{X} , and that the rejection mechanism assures that the buffer constraint is always met, namely:

$$\hat{X}^n(t) \in \mathcal{X}, \quad t \geq 0, \text{ a.s.} \tag{12}$$

For example, the case $\mathcal{X} = \{y \in \mathbb{R}_+^I : y_i \leq b_i, i \in \mathcal{I}\}$ corresponds to a system having a dedicated buffer, of size $b_i\sqrt{n}$, for each class, i . A single, shared buffer of size $b\sqrt{n}$ can be modeled by letting $\mathcal{X} = \{y \in \mathbb{R}_+^I : \sum_i y_i \leq b\}$. In any case, the actual (un-normalized) buffer size scales like \sqrt{n} . To meet the constraint (12), the control mechanism must reject some of the arrivals. In particular, consider a class- i arrival occurring at a time t when

$$(X^n(t-) + e^{(i)})/\sqrt{n} \notin \mathcal{X}. \tag{13}$$

This arrival has to be rejected so as to keep (12) valid. Physically, this situation represents buffers being full, with no available space to accommodate new arrivals. Such rejections, that occur when (13) holds, are often called *loss* in the literature. In our setting, admission/rejection decisions are controlled by the decision maker, and it is natural to refer to these as part of the rejection control process. We will refer to them as *forced rejections*, to distinguish them from rejections that occur when the buffers are not full (i.e., when (13) does not hold).

The process $U^n = (Z^n, B^n)$ is regarded a control, that is determined based on observations from the past (and present) events in the system. The precise definition is as follows.

Definition 2.1. (Admissible control, QCP) Fix $n \in \mathbb{N}$ and consider fixed processes (A^n, S^n) given by (2) and (4). A^n and S^n are called the **primitive processes**. A process $U^n = (Z^n, B^n)$, taking values in $\mathbb{R}_+^I \times \mathcal{B}$, having RCLL sample paths with the processes $Z_i^n, i \in \mathcal{I}$ having nondecreasing sample paths and given in the form (7), is said to be an **admissible control** for the n -th system if the following holds. Let the processes T^n, D^n, X^n be defined by the primitive and control processes, (A^n, S^n) and (Z^n, B^n) , via equations (8), (9) and (10), respectively. Then

- (Z^n, B^n) is adapted to the filtration $\{\mathcal{F}_t^n\}_{t \geq 0}$, where $\mathcal{F}_t^n = \sigma\{A_i^n(s), D_i^n(s), i \in \mathcal{I}, s \leq t\}$;
- One has a.s., that, for all $i \in \mathcal{I}$ and $t \geq 0$,

$$X_i^n(t) = 0 \quad \text{implies} \quad B_i^n(t) = 0. \tag{14}$$

An admissible control under which the scaled version \hat{X}^n of X^n satisfies (12) is said to **satisfy the buffer constraints**.

The first bullet above asserts that control decisions are based on the past arrival and departure events. The second bullet expresses the fact that jobs from a certain class can be processed only if there is at least one customer of that class in the system. We denote the class of all admissible controls U^n by $\tilde{\mathcal{U}}^n$, and the subset of those members of $\tilde{\mathcal{U}}^n$ satisfying the buffer constraints, by \mathcal{U}^n . Except for the last section of this paper, we will refer to processes in \mathcal{U}^n as merely admissible controls, for short. Note that the class \mathcal{U}^n depends on the processes A^n and S^n , but we consider these processes to be fixed.

Fix $\alpha > 0$, $h \in (0, \infty)^I$ and $r \in (0, \infty)^I$. For each $n \in \mathbb{N}$ consider the cost

$$J^n(U^n) = \mathbb{E} \left[\int_0^\infty e^{-\alpha t} [h \cdot \hat{X}^n(t) dt + r \cdot d\hat{Z}^n(t)] \right], \quad U^n = (Z^n, B^n) \in \tilde{\mathcal{U}}^n. \quad (15)$$

It will be assumed throughout that, for some $x_0 \in \mathcal{X}$,

$$\hat{X}^n(0) \rightarrow x_0, \quad \text{as } n \rightarrow \infty. \quad (16)$$

The QCP value is given by

$$V^n = \inf_{U^n \in \mathcal{U}^n} J^n(U^n). \quad (17)$$

We will be interested in the asymptotic behavior of V^n .

Denote by $\theta^n = (\theta_i^n)_{i \in \mathcal{I}}$ $\theta_i^n = 1/\mu_i^n$, and $\theta = (\theta_i)_{i \in \mathcal{I}}$, $\theta_i = 1/\mu_i$. The process $\theta^n \cdot X^n$, its normalized version $\theta^n \cdot \hat{X}^n$ and the formal limit of the latter, $\theta \cdot X$, will play an important role in reducing the dimensionality of the problem. These processes are often referred to as the *nominal workload* (eg., in [34]), but we will refer to them simply as *workload*.

2.2 The Brownian control problems

Using (6), (10) and the definition of the rescaled processes, a simple calculation shows that the following identity holds for $i \in \mathcal{I}$ and $t \geq 0$:

$$\hat{X}_i^n(t) = \hat{X}_i^n(0) + \hat{W}_i^n(t) + \hat{Y}_i^n(t) - \hat{Z}_i^n(t), \quad (18)$$

where, denoting $m_i = \hat{\lambda}_i - \rho_i \hat{\mu}_i$,

$$m_i^n = \frac{\lambda_i^n - \rho_i \mu_i^n}{\sqrt{n}} = m_i + o(1), \quad (19)$$

$$\hat{W}_i^n(t) = \hat{A}_i^n(t) - \hat{S}_i^n(T_i^n(t)) + m_i^n t, \quad (20)$$

and

$$\hat{Y}_i^n(t) = \frac{\mu_i^n}{\sqrt{n}} (\rho_i t - T_i^n(t)). \quad (21)$$

Since $\sum_i \rho_i = 1$ and one always has $\sum_i B_i^n(t) \leq 1$, it follows that

$$\theta^n \cdot \hat{Y}^n \quad \text{is a nonnegative, nondecreasing process.} \quad (22)$$

We derive from (18)–(22) and (15)–(17) a control problem associated with diffusion by taking formal limits. Consider equation (18). The scaled initial conditions converge to x by (16). Next, the centered, rescaled renewal process \hat{A}_i^n [resp., \hat{S}_i^n] converges weakly to a BM starting from zero, with zero mean and diffusion coefficient $\sqrt{\lambda_i}C_{IA^i}$ [resp., $\sqrt{\mu_i}C_{ST^i}$] (see Section 17 of [13]). Heuristically, if the processes involved in (18) are to give rise to a limiting BCP then in particular \hat{Y}^n are order one as $n \rightarrow \infty$. Thus by (21) one has that $T^n(t)$ converge to ρt . Thus, taking into account the time change in the second term of (20), \hat{W}^n is to be replaced a BM starting from zero, with drift vector $m = (m_i)_{i \in \mathcal{I}}$ and diffusion matrix $\sigma = \text{diag}(\sigma_i)$, where

$$\sigma_i^2 := \lambda_i C_{IA^i}^2 + \mu_i C_{ST^i}^2 \rho_i = \lambda_i (C_{IA^i}^2 + C_{ST^i}^2).$$

Such a process will be called an (m, σ) -BM. Finally, \hat{Y}^n gives rise to a process Y for which $\theta \cdot Y$ is nonnegative and nondecreasing, whereas \hat{Z}^n to a process having nonnegative, nondecreasing components.

2.2.1 The BCP

Definition 2.2. (Admissible control, BCP) *An admissible control for the initial condition $x_0 \in \mathcal{X}$ is a filtered probability space $(\Omega', \mathcal{F}', \{\mathcal{F}'_t\}, \mathbb{P}')$ for which there exist an (m, σ) -BM, W , and a process $U = (Y, Z)$ taking values in $(\mathbb{R}_+^I)^2$, with RCLL sample paths, such that the following conditions hold:*

- W , Y and Z are adapted to $\{\mathcal{F}'_t\}$;

-

$$\text{For } 0 \leq s < t, \text{ the increment } W(t) - W(s) \text{ is independent of } \mathcal{F}'_s \text{ under } \mathbb{P}'; \quad (23)$$

-

$$\theta \cdot Y \text{ and } Z_i, i = 1, \dots, I, \text{ are nondecreasing}; \quad (24)$$

- With

$$X(t) = x_0 + W(t) + Y(t) - Z(t), \quad t \geq 0, \quad (25)$$

one has

$$X(t) \in \mathcal{X} \quad \text{for all } t, \mathbb{P}'\text{-a.s.} \quad (26)$$

We write $\mathcal{A}(x_0)$ for the class of admissible controls for the initial condition x_0 . When we write $(Y, Z) \in \mathcal{A}(x_0)$ it will be understood that these processes carry with them a filtered probability space and the processes W and X . Moreover, with a slight abuse of notation, we will write \mathbb{E} for the expectation corresponding to this probability space. For $(Y, Z) \in \mathcal{A}(x_0)$, let

$$J(x_0, Y, Z) = \mathbb{E} \left[\int_0^\infty e^{-\alpha t} [h \cdot X(t) dt + r \cdot dZ(t)] \right]. \quad (27)$$

The BCP is to find (Y, Z) that minimize $J(Y, Z)$ and achieve the value

$$V(x_0) = \inf_{(Y, Z) \in \mathcal{A}(x_0)} J(x_0, Y, Z). \quad (28)$$

2.2.2 The RBCP

The BCP is treated by reduction to a one-dimensional problem. This is obtained by multiplying equation (25) and the processes involved in it by θ . To introduce it, denote $\bar{x}_0 = \theta \cdot x_0$, $\bar{m} = \theta \cdot m$ and $\bar{\sigma}^2 = \sum \theta_i^2 \sigma_i^2$. Let

$$\mathbf{x} = \max\{\theta \cdot \xi : \xi \in \mathcal{X}\}. \quad (29)$$

Definition 2.3. (Admissible control, RBCP) *An admissible control for the initial condition $\bar{x}_0 \in [0, \mathbf{x}]$ is a filtered probability space $(\Omega', \mathcal{F}', \{\mathcal{F}'_t\}, \mathbb{P}')$ for which there exist an $(\bar{m}, \bar{\sigma})$ -BM, \bar{W} , and a process $\bar{U} = (\bar{Y}, \bar{Z})$ taking values in \mathbb{R}_+^2 , with RCLL sample paths, such that the following conditions hold:*

- \bar{W} , \bar{Y} and \bar{Z} are adapted to $\{\mathcal{F}'_t\}$;
- For $0 \leq s < t$, the increment $\bar{W}(t) - \bar{W}(s)$ is independent of \mathcal{F}'_s under \mathbb{P}' ;
-

$$\bar{Y} \text{ and } \bar{Z} \text{ are nondecreasing;} \quad (30)$$

- With

$$\bar{X}(t) = \bar{x}_0 + \bar{W}(t) + \bar{Y}(t) - \bar{Z}(t), \quad t \geq 0, \quad (31)$$

one has

$$\bar{X}(t) \in [0, \mathbf{x}] \quad \text{for all } t, \mathbb{P}'\text{-a.s.} \quad (32)$$

We write $\bar{\mathcal{A}}(\bar{x}_0)$ for the class of admissible controls for the initial condition \bar{x}_0 . Given $(\bar{Y}, \bar{Z}) \in \bar{\mathcal{A}}(\bar{x}_0)$, let

$$\bar{J}(\bar{x}_0, \bar{Y}, \bar{Z}) = \mathbb{E} \left[\int_0^\infty e^{-\alpha t} [\bar{h}(\bar{X}(t)) dt + \bar{r} d\bar{Z}(t)] \right], \quad (33)$$

where

$$\begin{aligned} \bar{h}(w) &= \min\{h \cdot \xi : \xi \in \mathcal{X}, \theta \cdot \xi = w\}, \quad w \in [0, \mathbf{x}], \\ \bar{r} &= \min\{r \cdot z : z \in \mathbb{R}_+^I, \theta \cdot z = 1\}. \end{aligned}$$

Note that \bar{h} is convex by convexity of the set \mathcal{X} (in case when \mathcal{X} is polyhedral, \bar{h} is also piecewise linear). Note also that as members of $(0, \infty)^I$, θ and h cannot be orthogonal, thus $\bar{h}(w) > 0$ for any $w > 0$. Since $\bar{h}(0) = 0$, it follows that \bar{h} is strictly increasing. Let

$$\bar{V}(\bar{x}_0) = \inf_{(\bar{Y}, \bar{Z}) \in \bar{\mathcal{A}}(\bar{x}_0)} \bar{J}(\bar{x}_0, \bar{Y}, \bar{Z}).$$

Toward relating the two problems, we will need the following additional definitions. First, the extremal points of the set $\{z \in \mathbb{R}_+^I : \theta \cdot z = 1\}$ are precisely $\theta_i^{-1} e^{(i)}$, namely $\mu_i e^{(i)}$, $i \in \mathcal{I}$. Hence there exists (at least one) i^* such that $\zeta^* = \mu_{i^*} e^{(i^*)}$ satisfies

$$\zeta^* \in \arg \min_z \{r \cdot z : z \in \mathbb{R}_+^I, \theta \cdot z = 1\}.$$

Fix such i^* and the corresponding ζ^* . Note that i^* can alternatively be characterized via

$$r_{i^*} \mu_{i^*} = \min_i r_i \mu_i. \quad (34)$$

Next, let $\gamma : [0, \mathbf{x}] \rightarrow \mathcal{X}$ be Borel measurable, satisfying

$$\gamma(w) \in \arg \min_{\xi} \{h \cdot \xi : \xi \in \mathcal{X}, \theta \cdot \xi = w\}, \quad w \in [0, \mathbf{x}]. \quad (35)$$

(For the existence of a measurable selection see Corollary 10.3 in the appendix of [22]). Note that, by definition, $\gamma(w) \in \mathcal{X}$, $\theta \cdot \gamma(w) = w$, and $h \cdot \gamma(w) = \bar{h}(w) \leq h \cdot \xi$ for every $\xi \in \mathcal{X}$ for which $\theta \cdot \xi = w$. The relation between the problems is as follows.

Proposition 2.1. *Let $x_0 \in \mathcal{X}$ and $\bar{x}_0 = \theta \cdot x_0$.*

- i. Given an admissible control $(\Omega', \mathcal{F}', \{\mathcal{F}'_t\}, \mathbb{P}', W, Y, Z)$ for x for the (multidimensional) BCP, define $(\bar{W}, \bar{X}, \bar{Y}, \bar{Z})$ by $(\theta \cdot W, \theta \cdot X, \theta \cdot Y, \theta \cdot Z)$. Then $(\bar{Y}, \bar{Z}) \in \bar{\mathcal{A}}(\bar{x}_0)$ and $\bar{J}(\bar{x}_0, \bar{Y}, \bar{Z}) \leq J(x_0, Y, Z)$.*
- ii. Conversely, let an admissible control $(\Omega', \mathcal{F}', \{\mathcal{F}'_t\}, \mathbb{P}', \bar{W}, \bar{Y}, \bar{Z})$ for \bar{x}_0 for the RBCP be given, and assume the probability space supports an (m, σ) -BM W . Assume W is $\{\mathcal{F}'_t\}$ -adapted and satisfies $\theta \cdot W = \bar{W}$ and (23). Construct (X, Y, Z) by*

$$X(t) = \gamma(\bar{X}(t)), \quad Z(t) = \zeta^* \bar{Z}(t), \quad (36)$$

$$Y(t) = X(t) - x_0 - W(t) + Z(t). \quad (37)$$

Then $(Y, Z) \in \mathcal{A}(x_0)$, and $J(x_0, Y, Z) \leq \bar{J}(\bar{x}_0, \bar{Y}, \bar{Z})$.

iii. $V(x_0) = \bar{V}(\bar{x}_0)$.

Proof. i. We verify that Definition 2.3 is satisfied by $(\bar{W}, \bar{X}, \bar{Y}, \bar{Z})$. The first three bullets in that definition are straightforward. Equation (31) follows from (25), while (32) from (26).

Now, by definition of \bar{h} ,

$$h \cdot X(t) \geq \bar{h}(\theta \cdot X(t)), \quad (38)$$

and by definition of \bar{r} ,

$$\int_0^\infty e^{-\alpha t} r \cdot dZ(t) \geq \int_0^\infty e^{-\alpha t} \bar{r} d(\theta \cdot Z(t)). \quad (39)$$

Therefore

$$J(x_0, Y, Z) \geq \bar{J}(\bar{x}_0, \bar{Y}, \bar{Z}).$$

ii. We show that $(Y, Z) \in \mathcal{A}(x_0)$ by verifying that Definition 2.2 is satisfied. The adaptedness follows by the assumption on W and the construction of X, Y and Z . Property (23) holds by assumption. By construction, (X, Y, Z) satisfy (25). Property (26) holds because, by definition, $\gamma(w) \in \mathcal{X}$ for all $w \in [0, \mathbf{x}]$. Z_i are nonnegative, nondecreasing because so is \bar{Z} , and $\zeta_i^* \geq 0$. Moreover,

$$\theta \cdot Y(t) = \theta \cdot X(t) - \theta \cdot x_0 - \theta \cdot W(t) + \theta \cdot Z(t) = \bar{Y}(t).$$

Hence $\theta \cdot Y$ is nonnegative, nondecreasing. As a result, $(Y, Z) \in \mathcal{A}(x_0)$.

Next, note that

$$h \cdot X(t) = \bar{h}(\bar{X}(t)), \quad r \cdot dZ(t) = \bar{r} d\bar{Z}(t).$$

Therefore

$$J(x_0, Y, Z) = \bar{J}(\bar{x}_0, \bar{Y}, \bar{Z}). \quad (40)$$

iii. The last assertion will follow from the first two once we show that, in (ii), one can always find W with the stated properties. This is possible by supplementing the one-dimensional BM \bar{W} with an $(I - 1)$ -dimensional BM, independent of \bar{W} , and augmenting the probability space accordingly. Specifically, if \bar{W} is a (one-dimensional) $(\bar{m}, \bar{\sigma})$ -BM w.r.t. a filtration $\{\bar{\mathcal{F}}_t\}_{t \geq 0}$ and \hat{W} is a standard $(I - 1)$ -dimensional BM independent of \bar{W} then it is not hard to see that an $I \times (I - 1)$ matrix \hat{A} and I -dimensional vectors \bar{A} and a can be found so that the I -dimensional process $W(t) = \hat{A}\hat{W}(t) + \bar{A}\bar{W}(t) + at$ is an (m, σ) -BM and one has $\theta \cdot W(t) = \bar{W}(t)$, $t \geq 0$. Letting

$$\mathcal{F}'_t = \bar{\mathcal{F}}_t \vee \sigma\{\hat{W}(s) : s \in [0, t]\}, \quad t \geq 0,$$

then gives a filtration with which all conditions of an admissible control for the BCP are satisfied. \square

2.2.3 The Harrison-Taksar free boundary problem

The solution to the one-dimensional problem has been studied by Harrison and Taksar [28] via the Bellman equation. They showed that the function \bar{V} is $C^2[0, \mathbf{x}]$ and solves the equation

$$\begin{cases} \left[\frac{1}{2} \bar{\sigma}^2 f'' + \bar{m} f' - \alpha f + \bar{h} \right] \wedge f' \wedge [\bar{r} - f'] = 0, & \text{in } (0, \mathbf{x}), \\ f'(0) = 0, \quad f'(\mathbf{x}) = \bar{r}. \end{cases} \quad (41)$$

It follows from their work that an optimal control is one under which the process \bar{X} is a RBM on a certain subinterval of $[0, \mathbf{x}]$. We will consider a RBM as a path transformation of a BM by a *Skorohod map*, a map that will later be used in a wider context. To introduce this map, let $a > 0$. The Skorohod map on the interval $[a, b]$, denoted by $\Gamma_{[a, b]}$, is map $D([0, \infty) : \mathbb{R}) \rightarrow D([0, \infty) : \mathbb{R})^3$. It is characterized as the solution map $\psi \rightarrow (\varphi, \eta_1, \eta_2)$ to the so called *Skorohod Problem*, namely the problem of finding, for a given ψ , a triplet $(\varphi, \eta_1, \eta_2)$, such that

$$\varphi = \psi + \eta_1 - \eta_2, \quad \varphi(t) \in [a, b] \text{ for all } t, \quad (42)$$

$$\eta_i \text{ are nonnegative and nondecreasing, } \eta_i(0-) = 0, \text{ and } \int_{[0, \infty)} 1_{(a, b]}(\varphi) d\eta_1 = \int_{[0, \infty)} 1_{[a, b)}(\varphi) d\eta_2 = 0. \quad (43)$$

By writing $\eta_i(0-) = 0$ we adopt the convention that $\eta_i(0) > 0$ is regarded a jump at zero. This convention, in conjunction with $\int_{[0, \infty)} 1_{(a, b]}(\varphi) d\eta_1 = 0$ [resp., $\int_{[0, \infty)} 1_{[a, b)}(\varphi) d\eta_2 = 0$], means that if $\psi(0) < a$ [resp., $\psi(0) > b$] then $\varphi(0) = a$ [resp., b]. If, however, $\psi(0) \in [a, b]$ then $\varphi(0) = \psi(0)$, and η_i have no jump at zero.

See [32] for existence and uniqueness of solutions, and continuity and further properties of the map. In particular, it is well-known that $\Gamma_{[a, b]}$ is continuous in the uniformly-on-compacts topology.

We now go back to the RBCP. The following is mostly a result of [28].

Proposition 2.2. *The function \bar{V} is in $C^2[0, \mathbf{x}]$ and solves (41) uniquely among all $C^2[0, \mathbf{x}]$ functions.¹ Denote $\mathbf{x}^* = \inf\{y \in [0, \mathbf{x}] : \bar{V}'(z) = \bar{r} \text{ for } z \in [y, \mathbf{x}]\}$. Then $\mathbf{x}^* \in (0, \mathbf{x})$. Fix $\bar{x}_0 \in [0, \mathbf{x}]$. Let*

¹This uniqueness question was left open in [28], at the end of Section 6.

\bar{W} be an $(\bar{m}, \bar{\sigma})$ -BM and let \bar{X} , \bar{Y} and \bar{Z} be the corresponding RBM on $[0, \mathbf{x}^*]$ and boundary terms for 0 and \mathbf{x}^* , defined as

$$(\bar{X}, \bar{Y}, \bar{Z}) = I|_{[0, \mathbf{x}^*]}(\bar{x}_0 + \bar{W}). \quad (44)$$

Then (\bar{Y}, \bar{Z}) is optimal for $\bar{V}(\bar{x}_0)$, i.e., $\bar{J}(\bar{x}_0, \bar{Y}, \bar{Z}) = \bar{V}(\bar{x}_0)$.

Remark 2.1. Note that \bar{X} has the form $\bar{X} = \bar{x}_0 + \bar{W} + \bar{Y} - \bar{Z}$. Moreover, if $\bar{x}_0 > \mathbf{x}^*$ then \bar{X} is initially at \mathbf{x}^* ; in particular, $\bar{Z}(0) = (\bar{x}_0 - \mathbf{x}^*)^+$.

Proof. The fact that \bar{V} is C^2 and solves the equation is proved in [28], Proposition 6.6 and the discussion that follows. Uniqueness follows from the uniqueness of solutions in the viscosity sense, for a class of equations for which the above is a special case [7]. Let us explain how. It follows from the main result of [7] that uniqueness of viscosity solutions holds for (41) where the Neumann boundary condition (BC) is replaced by a state constraint BC (see [7] for the definitions of viscosity solutions and state constraint BC). As is well-known (and follows directly from the definition), any C^2 function satisfying equation (41) is also a viscosity solution in the interior $(0, \mathbf{x})$. As for the state constraint BC, it is easy to check (again following directly from the definition) that any smooth function satisfying the Neumann BC $f'(0) = 0$ and $f'(\mathbf{x}) = \bar{r}$, also satisfies the state constraint BC. This gives the uniqueness.

It is shown in [28] (see the discussion preceding (6.9) therein) that $f' = \bar{r}$ on $[y, \mathbf{x}]$, some $y \in [0, \mathbf{x})$. This shows $\mathbf{x}^* < \mathbf{x}$.

Next, it is shown in [28] that the control under which \bar{X} is a RBM on $[a, \mathbf{x}^*]$, for some $0 \leq a < \mathbf{x}^*$, is optimal. It remains to show that, in the case considered in this paper, $a = 0$. The argument relies on the fact that \bar{h} is strictly increasing, as shown in the discussion following (33).

Arguing by contradiction, assume $a > 0$. The interval $[a, \mathbf{x}^*]$ is independent of the initial condition, and so we are free to choose any \bar{x}_0 . Consider $\bar{x}_0 = 0$. Consider a BM \bar{W} and the process $\bar{X} = \bar{x}_0 + \bar{W} + \bar{Y} - \bar{Z}$ that initially has the value a and is given as a RBM on $[a, \mathbf{x}^*]$, driven by \bar{W} . By the result of [28] alluded to above, (\bar{Y}, \bar{Z}) is optimal, i.e., $\bar{J}(0, \bar{Y}, \bar{Z}) = \bar{V}(0)$. Let τ be the first hitting time of \bar{X} at $a + \varepsilon < \mathbf{x}^*$. Next, construct on the same probability space another triplet $(\tilde{X}, \tilde{Y}, \tilde{Z})$, where \tilde{X} behaves as a RBM on $[0, \mathbf{x}^*]$, driven by \bar{W} , up to the time τ , and starting at time τ agrees with \bar{X} (in particular, it has a jump at time τ). In other words, $(\tilde{X}, \tilde{Y}, \tilde{Z}) = (\bar{X} - a, \bar{Y} - a, 0)$ on $[0, \tau)$, and $(\tilde{X}, \tilde{Y}, \tilde{Z}) = (\bar{X}, \bar{Y}, \bar{Z})$ on $[\tau, \infty)$. Clearly, the cost incurred by $(\tilde{X}, \tilde{Y}, \tilde{Z})$ on $[\tau, \infty)$, namely

$$\int_{[\tau, \infty)} e^{-\alpha t} [\bar{h}(\tilde{X}(t))dt + \bar{r}d\tilde{Z}(t)]$$

is equal to that incurred by $(\bar{X}, \bar{Y}, \bar{Z})$ on that interval, while, owing to the strict monotonicity of \bar{h} and the positivity of τ ,

$$\int_{[0, \tau)} e^{-\alpha t} \bar{h}(\tilde{X}(t))dt < \int_{[0, \tau)} e^{-\alpha t} \bar{h}(\bar{X}(t))dt.$$

Note that no cost of the form $\bar{r}d\tilde{Z}$ is incurred during the time interval $[0, \tau)$. Taking expectations shows $\bar{J}(0, \tilde{Y}, \tilde{Z}) < \bar{J}(0, \bar{Y}, \bar{Z}) = \bar{V}(0)$, a contradiction. This shows $a = 0$. \square

As an immediate consequence of the above two results, we obtain an optimal control for the BCP.

Corollary 2.1. *Let $x_0 \in \mathcal{X}$ and W be an (m, σ) -BM. Denote $\bar{x}_0 = \theta \cdot x_0$ and $\bar{W} = \theta \cdot W$, and let \mathbf{x}^* be the free boundary point. Let $(\bar{X}, \bar{Y}, \bar{Z})$ be defined in terms of \bar{W} as in (44), and let (X, Y, Z) be defined in terms of (W, \bar{X}, \bar{Z}) as in (36)–(37). Then (Y, Z) is optimal for $V(x_0)$, namely $J(x_0, Y, Z) = V(x_0)$.*

2.3 Discussion

A brief description of the solution to the BCP is as follows. The workload process $\bar{X} = \theta \cdot X$ is given as a RBM on $[0, \mathbf{x}^*]$, where the free boundary point \mathbf{x}^* is dictated by the Bellman equation. The multidimensional queue length process X is recovered from \bar{X} by $X = \gamma(\bar{X})$. The multidimensional rejection process Z has only one nonzero component, namely the i^* -th component, which increases only when $\bar{X} \geq \mathbf{x}^*$.

This structure has an interpretation for the queueing model, that can be used to identify asymptotically optimal policies. Our main interest will be in the case of a rectangular domain, namely

$$\mathcal{X} = \{x \in \mathbb{R}^I : 0 \leq x_i \leq b_i, i \in \mathcal{I}\}, \quad (45)$$

for some fixed $b_i > 0$, representing a system where each class has a dedicated buffer (this will be our assumption in Section 4, although in Section 3 we allow general domains). In this case, the parameter \mathbf{x} associated with the RBCP (defined in (29)) is given by $\theta \cdot b$.

The BCP solution suggests that, in the queueing model, rejections should occur only when the scaled workload exceeds the level \mathbf{x}^* , and only from class i^* . Recall from (34) that this class is the class for which $r_i \mu_i$ is minimal. As explained in [34], i^* is the class for which the rejection penalty per unit of work is smallest.

Next, the relation

$$\hat{X}^n = \gamma(\theta \cdot \hat{X}^n) + o(1) \quad (46)$$

between the queue length and workload processes should hold. This is a requirement on the scheduling control. As mentioned in the introduction, when a critically loaded multiclass G/G/1 queue operates under fixed priority, the queue length of all classes but one is asymptotically zero in diffusion scale, the exception being the class with least priority [42]. This is a simple example of a scheduling policy that dictates a relation of the form (46), where here $\gamma(w) = (0, \dots, 0, w \mu_I)$. Relation (46) with a more complicated γ appears implicitly when applying the generalized $c\mu$ rule of [40]. In [10] and [8] the scheduling policies keep (46) where γ is a generic minimizing curve.

We can solve for the minimizing curve γ in the present setting, where \mathcal{X} takes the form (45). Equation (35) can in this case be written as

$$\gamma(w) \in \arg \min_x \{h \cdot x : 0 \leq x_i \leq b_i \text{ for all } i, \text{ and } \theta \cdot x = w\}, \quad w \in [0, \mathbf{x}].$$

Assume that the classes are labeled in such a way that

$$h_1 \mu_1 \geq h_2 \mu_2 \geq \dots \geq h_I \mu_I. \quad (47)$$




Figure 1: The curve γ (thick line) for the case of dedicated buffers in dimension $I = 2$ and 3 . As workload increases, starting from level zero, queue length I builds up until the corresponding buffer becomes full. Then buffer $I - 1$ starts to build up, and so on, until the rejection level $(\mathbf{x}^*, \gamma(\mathbf{x}^*))$ is reached. Rejections that occur starting at \mathbf{x}^* assure that this level is not exceeded.




Figure 2: The case of a single buffer shared by 2 classes. Rejection level may be reached before the buffer with least priority is full (left). For a higher rejection level, the curve continues along the boundary (right).

Given $w \in [0, \mathbf{x}]$ let (j, ξ) be the unique pair determined by the relation

$$w = \sum_{i=j+1}^I \theta_i b_i + \theta_j \xi, \quad j \in \{1, 2, \dots, I\}, \xi \in [0, b_j]. \quad (48)$$

An exception is the special case $w = \mathbf{x} = \theta \cdot b$, where one lets $j = 1$ and $\xi = b_1$. In other words, denote $\hat{b}_j = \sum_{i=j+1}^I \theta_i b_i$, $j \in \{0, \dots, I\}$ and note that $0 = \hat{b}_I < \hat{b}_{I-1} < \dots < \hat{b}_1 < \hat{b}_0 = \theta \cdot b = \mathbf{x}$. Then $j = j(w)$ is determined by

$$w \in [\hat{b}_j, \hat{b}_{j-1})$$

and $\xi = \xi(w) = (w - \hat{b}_j) / \theta_j$. Thus γ can be written explicitly as

$$\gamma(w) = \sum_{i=j+1}^I b_i e^{(i)} + \xi e^{(j)}. \quad (49)$$

Simple examples are depicted in Figure 1. While the usual use of the $c\mu$ rule is by assigning fixed priority, here the index shows up differently. When buffer I becomes full and workload is increased, a queue in buffer $I - 1$ starts building up, and so on. A policy that aims at achieving (46) is developed in Section 4. Examples for the case of a shared buffer are depicted in Figure 2. As shown in this figure, the case of two classes with a shared buffer leads to a triangular domain. In higher dimension one may think of one set of classes sharing one buffer, another set sharing another buffer etc., leading to more examples of non-rectangular domains. General domains are covered in this paper as far as the lower bound is concerned, but we only address AO controls for the case of rectangular domains.

Example 2.1. (Numerical solution of the BCP) *In this example we consider a specific three-dimensional BCP and provide its solution explicitly. The parameters are given in the following table:*

i	b_i	h_i	r_i	μ_i	λ_i	$h_i\mu_i$	$r_i\mu_i$
1	15	32.9	5.0	28.0	9.33	921.2	140
2	15	35.0	4.0	23.0	7.67	805	92
3	10	39.0	5.5	18.0	6.0	702	99

We further assume that $\hat{\lambda}_i = \hat{\mu}_i = 0$ for all i (so that $\bar{m} = 0$), that $\bar{\sigma}^2 = 0.1$, and take the discount parameter $\alpha = 10$. The resulting ordering of the $h_i\mu_i$ index is as in (47), namely $h_1\mu_1 > h_2\mu_2 > h_3\mu_3$. The ordering of $r_i\mu_i$ is such that class 2 is the most inexpensive as far as rejections are concerned, that is, $i^* = 2$. The Bellman equation takes the form

$$\begin{cases} [0.05f'' - 10f + \bar{h}] \wedge f' \wedge [92 - f'] = 0, & \text{in } (0, 1.74), \\ f'(0) = 0, & f'(1.74) = 92. \end{cases} \quad (50)$$

The function \bar{h} defined by

$$\bar{h}(w) = \min \left\{ \sum_{i=1}^3 h_i \xi_i : \xi \in \mathcal{X}, \sum_{i=1}^3 \theta_i \xi_i = w \right\}, \quad w \in [0, 1.74],$$

where $\mathcal{X} = [0, 15] \times [0, 15] \times [0, 10]$ and $\theta_i = \mu_i^{-1}$, is explicitly given by

$$\bar{h}(w) \approx \begin{cases} 18 \cdot 39 w & 0 \leq w \leq 0.56, \\ 390 + 23 \cdot 35 (w - 0.56) & 0.56 < w \leq 1.21, \\ 915 + 28 \cdot 32.9 (w - 1.21) & 1.21 < w \leq 1.74. \end{cases}$$

A numerical solution of the equation is shown in Figure 3 below.

The free boundary point in this case, found numerically, is the point $\mathbf{x}^* \approx 1.47$ at which $V' = \bar{r}$.




Figure 3: Graphs of V and V' . The free boundary point \mathbf{x}^* is found by seeking the smallest x for which $V'(x) = \bar{r}$ (Example 2.1).

The curve γ from (48)–(49) is given by

$$\gamma(w) \approx \begin{cases} \frac{w}{0.56} 10 [0, 0, 1] & 0 \leq w \leq 0.56, \\ 10 [0, 0, 1] + \frac{w - 0.56}{1.21 - 0.56} 15 [0, 1, 0] & 0.56 < w \leq 1.21, \\ 10 [0, 0, 1] + 15 [0, 1, 0] + \frac{w - 1.21}{1.74 - 1.21} 15 [1, 0, 0] & 1.21 < w \leq 1.47. \end{cases}$$

(We have not specified γ for values of w beyond the free boundary point 1.47). Note that the structure of this curve is of the form depicted in Figure 1 (right).

Example 2.2. (Numerical solution of the QCP) Here we present simulation results for the behavior of a two-class $M/M/1$ queue operating under the optimal policy. While for general service time and inter-arrival time distributions finding the optimal policy is hard, in the case of Poisson arrivals and exponential service times the problem has the form of a Markov decision process and one has access to the optimal policy by means of the corresponding Bellman equation on the discrete 2d grid. We have solved this equation numerically, computed the optimal policy based on the solution, and run a simulation for the behavior of the resulting queue length process. Figure 4 depicts histograms for the position of the two-dimensional queueing process, where gray levels encode the frequency of visits to each site in the state space (darker gray corresponds to more often visited sites). The histograms are depicted for an increasing value of the heavy traffic parameter. The results clearly indicate that the behavior becomes closer and closer to that given by the limit curve of the form depicted in Figure 1 (left).

While this numerical analysis is related to our results, note carefully that the relation is indirect: the simulation runs demonstrate the behavior under the optimal policy, whereas our results




Figure 4: Histograms for queue length process under the optimal policy for an increasing value of the heavy traffic parameter. The buffer sizes are 15×15 , 30×30 and 50×50 (Example 2.2).




Figure 5: Ratio between the (simulated) cost under the proposed policy and the (computed) optimal cost as a function of \sqrt{n} . The graph shows values for $\sqrt{n} = 3, 5, 10$ and 20 . The corresponding buffer sizes are given by $5\sqrt{n}$, namely 15×15 , 25×25 , 50×50 and 100×100 , respectively.

address the asymptotics under a sub-optimal (but AO) policy. The two are related in that both show convergence to the limit behavior identified by the BCP solution.

Finally, we have also simulated the performance of the sub-optimal policy that we propose. The graph in Figure 5 shows the ratio between the cost under the proposed policy and the optimal cost for different values of n .

3 A general lower bound

Recall that V^n is defined for the specific initial condition $\hat{X}^n(0)$, and that by (16), $\hat{X}^n(0) \rightarrow x_0$ as $n \rightarrow \infty$. The main result of this section asserts that the performance of any sequence of policies for the queueing model is asymptotically bounded below by the BCP value function.

Theorem 3.1. $\underline{V} := \liminf_{n \rightarrow \infty} V^n \geq V(x_0)$.

With an eye toward the last section, we will, in fact, prove a slightly stronger result. Instead of assuming the hard constraint (12), that is a part of the definition of ‘admissible controls satisfying the buffer constraint’, we will assume throughout this section the following weaker condition.

$$\begin{aligned} &\text{For every open set } \tilde{\mathcal{X}} \subset \mathbb{R}^I \text{ with } \mathcal{X} \subset \tilde{\mathcal{X}}, \text{ and every } T > 0, \\ &\mathbb{P}(\hat{X}^n(t) \in \tilde{\mathcal{X}} \text{ for all } t \in [0, T]) \rightarrow 1 \text{ as } n \rightarrow \infty. \end{aligned} \quad (51)$$

Denote by \mathcal{J} the operator

$$\mathcal{J}\varphi = \int_0^\cdot \varphi(t) dt,$$

for locally integrable functions φ .

Lemma 3.1. For $U^n \in \tilde{\mathcal{U}}^n$,

$$J^n(U^n) = \tilde{J}^n(U^n) := \mathbb{E} \left[\int_0^\infty e^{-\alpha t} [\alpha h \cdot \mathcal{J}\hat{X}^n(t) + \alpha^2 r \cdot \mathcal{J}\hat{Z}^n(t)] dt \right]. \quad (52)$$

Proof. The identity will follow from integration by parts once we show that the three terms $e^{-\alpha t} \mathcal{J}\hat{X}^n(t)$, $e^{-\alpha t} \hat{Z}^n(t)$ and $e^{-\alpha t} \mathcal{J}\hat{Z}^n(t)$ converge to zero a.s. as $t \rightarrow \infty$. Note by (10) that

$$\hat{X}_i^n(t) + \hat{Z}_i^n(t) = n^{-1/2} [X_i^n(t) + Z_i^n(t)] \leq n^{-1/2} [X_i^n(0) + A_i^n(t)].$$

As a renewal process with finite expectation, A_i^n satisfies a law of large numbers in the sense that $A_i^n(t)/t$ converges a.s. as $t \rightarrow \infty$. Thus the three terms alluded to above converge to zero a.s., and the identity follows. \square

Before stating the following lemma we introduce some additional notation. Let $\mathcal{A}^P(x_0)$ denote the class of controls for the BCP, defined as in Definition 2.2, except that instead of having RCLL paths, the processes are only assumed to be progressively measurable. More precisely, an element of $\mathcal{A}^P(x_0)$ is a filtered probability space $(\Omega', \mathcal{F}', \{\mathcal{F}'_t\}, \mathbb{P}')$ with an (m, σ) -BM, W , and a progressively measurable process (Y, Z) taking values in $(R_+^I)^2$, such that W is adapted, $W(t) - W(s)$ is independent of \mathcal{F}'_s ($0 \leq s < t$), and, on an event having full \mathbb{P}' -measure one has: $\theta \cdot Y$ is a.e. equal to a process with nondecreasing sample paths; the same holds for each of the processes Z_i , $i = 1, \dots, I$; and, with $X(t) = x_0 + W(t) + Y(t) - Z(t)$,

$$X(t) \in \mathcal{X} \quad \text{for a.e. } t.$$

Note that $\mathcal{A}(x_0) \subset \mathcal{A}^P(x_0)$.

The purpose of introducing this extended class of controls is as follows. The technique employed in the proof of Theorem 3.1 below is based on tightness of the processes $\mathcal{J}\hat{X}^n$, $\mathcal{J}\hat{Y}^n$ and $\mathcal{J}\hat{Z}^n$ rather than \hat{X}^n , \hat{Y}^n and \hat{Z}^n . It is established that the limits of these processes have Lipschitz continuous sample paths, and as a result they are a.e. differentiable. In order to connect these limits to the BCP one needs to construct from them an admissible control for the latter, but since the derivatives of Lipschitz functions need not be RCLL, the class of controls $\mathcal{A}(x_0)$ is too small for this purpose.

Using instead the class $\mathcal{A}^P(x_0)$ is possible thanks to a result from [20] (see below) that shows that progressively measurable a.e. derivatives always exist.

The following lemma shows that working with the extended class of controls does not vary the value function.

Lemma 3.2. *Let*

$$\tilde{J}(x_0, Y, Z) = \mathbb{E} \left[\int_0^\infty e^{-\alpha t} [\alpha h \cdot \mathcal{J}X(t) + \alpha^2 r \cdot \mathcal{J}Z(t)] dt \right]$$

and $V^P(x_0) = \inf_{\mathcal{A}^P(x_0)} \tilde{J}(x_0, Y, Z)$ (compare with the definitions (27), (28) of J and V). Then $V^P = V$.

Proof. Given x_0 , consider the specific control that is optimal for $V(x_0)$, namely (X, Y, Z) given in Proposition 2.1(ii), where $(\bar{X}, \bar{Y}, \bar{Z})$ is the RBM on $[0, \mathbf{x}^*]$. In particular, $Z(t) = \zeta^* \bar{Z}(t)$, where \bar{Z} is one of the boundary terms of a RBM. It is well known that $e^{-\alpha t}(\bar{Z}(t) + \mathcal{J}\bar{Z}(t)) \rightarrow 0$ a.s., as $t \rightarrow \infty$. As a result, a similar statement holds for $e^{-\alpha t}(Z_i(t) + \mathcal{J}Z_i(t))$, for each $i \in \mathcal{I}$. Using integration by parts, this shows that $V(x_0) = J(x_0, Y, Z) = \tilde{J}(x_0, Y, Z) \geq V^P(x_0)$.

Next, let $\varepsilon > 0$ and consider an ε -optimal control for $V^P(x_0)$, again denoted by (X, Y, Z) . Fix $T > 0$. Construct processes $(\tilde{X}, \tilde{Y}, \tilde{Z})$ that are identical to (X, Y, Z) on $[0, T]$. As for the time interval $[T, \infty)$, let \tilde{X} be a RBM on $[0, \mathbf{x}^*]$ starting from $\tilde{X}(T) = 0$, and let $(\tilde{X}, \tilde{Y}, \tilde{Z})$ be constructed from this RBM in the same fashion that (X, Y, Z) are constructed from \bar{X} in the first part of the proof. In particular, \tilde{Z} satisfies $e^{-\alpha t}(\tilde{Z}_i(t) + \mathcal{J}\tilde{Z}_i(t)) \rightarrow 0$, for each $i \in \mathcal{I}$ (and it may have a jump at T). By construction, (\tilde{Y}, \tilde{Z}) is progressively measurable. The constructed processes thus form an element of $\mathcal{A}^P(x_0)$, and owing to the above tail condition, using integration by parts, $J(x_0, \tilde{Y}, \tilde{Z}) = \tilde{J}(x_0, \tilde{Y}, \tilde{Z})$. Now,

$$\begin{aligned} \tilde{J}(x_0, \tilde{Y}, \tilde{Z}) - \tilde{J}(x_0, Y, Z) &= \mathbb{E} \left[\int_T^\infty e^{-\alpha t} [\alpha h \cdot (\mathcal{J}\tilde{X}(t) - \mathcal{J}X(t)) + \alpha^2 r \cdot (\mathcal{J}\tilde{Z}(t) - \mathcal{J}Z(t))] dt \right] \\ &\leq c \int_T^\infty t e^{-\alpha t} dt + \alpha^2 \int_T^\infty e^{-\alpha t} \int_T^t \mathbb{E}[r \cdot (\tilde{Z}(s) - \tilde{Z}(T))] ds dt \\ &\quad + \mathbb{E}[r \cdot \tilde{Z}(T)] \alpha^2 \int_T^\infty e^{-\alpha t} \int_T^t ds dt, \end{aligned}$$

using the equality $\tilde{Z} = Z$ on $[0, T]$. Since on $[T, \infty)$, $\tilde{Z} - \tilde{Z}(T)$ is the boundary term of an RBM on a fixed interval, it is a standard fact that the second term in the above display converges to zero as $T \rightarrow \infty$. As for the last term, since $\tilde{J}(x_0, \tilde{Y}, \tilde{Z}) < \infty$, one has $\mathbb{E} \int_T^\infty e^{-\alpha t} r \cdot \mathcal{J}\tilde{Z}(t) dt \rightarrow 0$ as $T \rightarrow \infty$, thus using monotonicity of $r \cdot \tilde{Z}$,

$$e^{-\alpha(T+2)} \mathbb{E}[r \cdot \tilde{Z}(T)] \leq e^{-\alpha(T+2)} \mathbb{E}[r \cdot \mathcal{J}\tilde{Z}(T+1)] \leq \mathbb{E} \int_{T+1}^{T+2} e^{-\alpha t} r \cdot \mathcal{J}\tilde{Z}(t) dt \rightarrow 0,$$

as $T \rightarrow \infty$. This shows

$$J(x_0, \tilde{Y}, \tilde{Z}) = \tilde{J}(x_0, \tilde{Y}, \tilde{Z}) \leq V^P(x_0) + \varepsilon + a(T),$$

where $a(T) \rightarrow 0$ as $T \rightarrow \infty$. Taking $T \rightarrow \infty$ shows $J(x_0, \tilde{Y}, \tilde{Z}) \leq V^P(x_0) + \varepsilon$. Thus to complete the proof, it suffices to show that $J(x_0, \tilde{Y}, \tilde{Z}) \geq V(x_0)$. This is not immediate, because (\tilde{Y}, \tilde{Z}) is an

element of $\mathcal{A}^P(x_0)$ whereas V is defined with the smaller class $\mathcal{A}(x_0)$. We will argue by passing to the one-dimensional problem. To this end, note that (38) and (39) are valid for the progressively measurable processes, thus

$$J(x_0, \tilde{Y}, \tilde{Z}) \geq \mathbb{E} \left[\int_0^\infty e^{-\alpha t} [\bar{h}(\theta \cdot \tilde{X}(t)) dt + \bar{r} d(\theta \cdot \tilde{Z}(t))] \right].$$

Now, the processes $\theta \cdot \tilde{Y}$ and $\theta \cdot \tilde{Z}$ are pathwise nondecreasing, due to the definition of $\mathcal{A}^P(x_0)$. Hence, if we define $\hat{Y}(t) = \lim_{s \downarrow t} \theta \cdot Y(s)$, $\hat{Z}(t) = \lim_{s \downarrow t} \theta \cdot Z(s)$ and $\hat{X} = \theta \cdot x_0 + \theta \cdot W + \hat{Y} - \hat{Z}$, then \hat{X} , \hat{Y} and \hat{Z} are RCLL. Moreover, they satisfy all assumptions of Definition 2.3, with $\bar{x}_0 = \theta \cdot x_0$ and $\bar{W} = \theta \cdot W$. As a result, they are in $\bar{\mathcal{A}}(\bar{x}_0)$, and so

$$J(x_0, \tilde{Y}, \tilde{Z}) \geq \bar{J}(\bar{x}_0, \hat{Y}, \hat{Z}) \geq \bar{V}(\bar{x}_0).$$

By Proposition 2.1, $\bar{V}(\bar{x}_0) = V(x_0)$. We have thus shown that $V(x_0) \leq V^P(x_0) + \varepsilon$, and the result follows on taking $\varepsilon \rightarrow 0$. \square

In the proof below and in the next section we will use the following characterization of C -tightness for processes with sample paths in $\mathbb{D}_{\mathbb{R}}$ (see Proposition VI.3.26 of [31]): *C -tightness of $\{X^N\}$, $N \in \mathbb{N}$ is equivalent to*

$$\text{The sequence of random variables } \|X^N\|_T \text{ is tight for every fixed } T < \infty, \text{ and} \quad (53)$$

For every $T < \infty$, $\varepsilon > 0$ and $\eta > 0$ there exist N_0 and $\theta > 0$ such that

$$N \geq N_0 \text{ implies } P(\bar{w}_T(X^N, \theta) > \eta) < \varepsilon, \quad (54)$$

where \bar{w} is defined in (1).

Proof of Theorem 3.1. The structure of the proof is as follows. We invoke Lemma 3.1 that allows us to work with the cost associated with the integrated version of the processes. We establish C -tightness of the integrated processes; more precisely, of the sequence $(\hat{W}^n, \mathcal{J}\hat{X}^n, \mathcal{J}\hat{Y}^n, \mathcal{J}\hat{Z}^n)$. The rest of the proof is devoted to showing that any subsequential limit of this sequence gives rise to control within the extended class \mathcal{A}^P , where the justification to work with the extended class is provided by Lemma 3.2.

We thus will rely on Lemma 3.1 and work with \tilde{J}^n . Using (17) and Lemma 3.1, $V^n = \inf_{\mathcal{U}^n} \tilde{J}^n(U^n)$. Fix a subsequence $\{n'\}$ along which $\lim \tilde{J}^{n'}(U^{n'}) = \underline{V}$, and relabel it as $\{n\}$. Assume, without loss of generality, that $\tilde{J}^n(U^n) < V(x_0) + 1$ for all n . Then $\tilde{J}^n(U^n)$ is bounded, and so is $J^n(U^n)$, and therefore, for every $T < \infty$,

$$e^{-\alpha T} \mathbb{E}[r \cdot \hat{Z}^n(T)] \leq \mathbb{E} \int_0^T r \cdot d\hat{Z}^n(t) \leq V(x_0) + 1.$$

This shows that $\|\hat{Z}^n(T)\|$, $n \in \mathbb{N}$, is tight as a sequence of r.v.s, for each T .

Recall that \hat{A}^n and \hat{S}^n converge u.o.c. to BMs, and note by (8) that $T_i^n(t) \leq t$ for every t . Using this and equations (19) and (20) shows that the sequence of processes \hat{W}^n is C -tight.

Given T , using the monotonicity of $\hat{Z}_i^n(\cdot)$, the Lipschitz constant of $\mathcal{J}\hat{Z}_i^n|_{[0,T]}$ is bounded by $\|\hat{Z}^n(T)\|$. Thus, using the characterization (53)–(54), the tightness of $\hat{Z}^n(T)$ for each T implies

that $\mathcal{J}\hat{Z}^n$ is a C -tight sequence of processes. The condition (51) implies that, for every T , $\|\hat{X}^n\|_T$, $n \in \mathbb{N}$, is a tight sequence of random variables. As a result, by (53)–(54), the sequence $\mathcal{J}\hat{X}^n$ is also C -tight. Next, by (18),

$$\|\hat{Y}^n(t)\| \leq \|\hat{X}^n(t)\| + \|\hat{W}^n(t)\| + \|\hat{Z}^n(t)\|. \quad (55)$$

It follows from this discussion that, for each T ,

$$L_n(T) := \|\hat{X}^n\|_T \vee \|\hat{Y}^n\|_T \vee \|\hat{Z}^n\|_T$$

is a tight sequence of r.v.s, and that $(\mathcal{J}\hat{X}^n, \mathcal{J}\hat{Y}^n, \mathcal{J}\hat{Z}^n)$ is C -tight, with bound $L_n(T)$ on the Lipschitz constant over the interval $[0, T]$. Since $L_n(T)$ are tight for each T , any weak limit point of the C -tight sequence is a process having locally Lipschitz paths a.s.

Next, since for each T , the sequence $\|\hat{Y}^n\|_T$ is tight, we have by (21) and the fact $\mu_i^n/\sqrt{n} \rightarrow \infty$, that T_i^n converge u.o.c. to \bar{T}_i where $\bar{T}_i(t) = \rho_i t$. By (20), using a lemma regarding random change of time [13], p. 151, it follows that $\hat{W}^n \Rightarrow W$, where we recall that W is an (m, σ) -BM.

By tightness of $(\hat{W}^n, \mathcal{J}\hat{X}^n, \mathcal{J}\hat{Y}^n, \mathcal{J}\hat{Z}^n)$, there exists a convergent subsequence. Denote its limit by (W, IX, IY, IZ) . Note that the last three terms have Lipschitz sample paths. By an argument as in section IV.17 of [20], they possess a.e. derivatives that are progressively measurable w.r.t. the filtration $\mathcal{F}'_t = \sigma\{W(s), IX(s), IY(s), IZ(s) : s \leq t\}$. For concreteness, let $\partial^- f$, for a Lipschitz $f : [0, \infty) \rightarrow \mathbb{R}$, be defined by $\partial^- f(0) = 0$ and $\partial^- f(t) = \liminf_{s \uparrow t} (t-s)^{-1}(f(t) - f(s))$, $t > 0$. Define pathwise $X = (X_i)$, $Y = (Y_i)$ and $Z = (Z_i)$ as $X_i = \partial^- IX_i$, $Y_i = \partial^- IY_i$ and $Z_i = \partial^- IZ_i$. Then (X, Y, Z) are progressively measurable, and $\mathcal{J}X = IX$. We will show below that these processes along with the filtration $\{\mathcal{F}'_t\}$ form an element of the class $\mathcal{A}^P(x_0)$. Consequently, using Lemma 3.1 and Fatou's lemma for the subsequence under consideration,

$$\underline{V} = \liminf \tilde{J}^n(U^n) \geq \tilde{J}(x_0, Y, Z) \geq \inf_{\mathcal{A}^P(x_0)} \tilde{J}(x_0, \cdot, \cdot) = V^P(x_0) = V(x_0),$$

where in the last equality we used Lemma 3.2.

It thus remains to show that the progressively measurable processes we have constructed form an element of the class $\mathcal{A}^P(x_0)$. To show (23), we borrow a few lines from the proof of Lemma 6 of [9]. Fix $0 \leq s \leq t < t + u$. Let $\alpha^n = (\hat{W}^n(s), \mathcal{J}\hat{X}^n(s), \mathcal{J}\hat{Y}^n(s), \mathcal{J}\hat{Z}^n(s))$ and $\alpha = (W(s), IX(s), IY(s), IZ(s))$. For $i \in \mathcal{I}$ let t_i^n [resp., τ_i^n] denote the renewal epoch of A_i^n [resp., S_i^n] following t [resp., $T_i^n(t)$]. That is,

$$t_i^n = \inf\{t' \geq t : A_i^n(t') > A_i^n(t)\}, \quad \tau_i^n = \inf\{t' \geq T_i^n(t) : S_i^n(t') > S_i^n(T_i^n(t))\}.$$

Let $\beta^n = (\beta_i^n)_{i \in \mathcal{I}}$ be defined by

$$\beta_i^n = (A_i^n(t_i^n + u) - A_i^n(t_i^n), S_i^n(\tau_i^n + \rho_i u) - S_i^n(\tau_i^n)).$$

Then α^n and β^n are mutually independent. As a result, α^n and $\gamma^n = (\gamma_i^n)_{i \in \mathcal{I}}$ are mutually independent, where

$$\gamma_i^n = \hat{A}_i^n(t_i^n + u) - \hat{S}_i^n(\tau_i^n + u) - \hat{A}_i^n(t_i^n) + \hat{S}_i^n(\tau_i^n) + m_i^n u.$$

Recall the definition (20) of \hat{W}^n . We have $t_i^n \Rightarrow t$, and $T_i^n(t) \Rightarrow \bar{T}(t)$ by which $\tau_i^n \Rightarrow \rho_i t$. As a result, $\hat{W}_i^n(t + u) - \hat{W}_i^n(t) - \gamma^n \Rightarrow 0$. This shows that α and $W(t + u) - W(t)$ are mutually

independent. Since $u > 0$ and $s \leq t$ are arbitrary, an application of Theorem 1.4.2 of [21] shows that all increments $W(t+u) - W(t)$ and \mathcal{F}'_s are independent.

Let $\mathcal{X}_\delta = \{x \in \mathbb{R}^I : \text{dist}(x, \mathcal{X}) < \delta\}$, $\delta > 0$. Condition (51) implies that for every $s < t$ and $\delta > 0$, $(t-s)^{-1}(\mathcal{J}\hat{X}^n(t) - \mathcal{J}\hat{X}^n(s)) \in \mathcal{X}_\delta$ occurs with probability tending to 1 as $n \rightarrow \infty$. As a result, $(t-s)^{-1}(IX(t) - IX(s)) \in \overline{\mathcal{X}_\delta}$ a.s. Since \mathcal{X} is closed and convex, the intersection of $\overline{\mathcal{X}_\delta}$ over $\delta > 0$ gives \mathcal{X} , so $(t-s)^{-1}(IX(t) - IX(s)) \in \mathcal{X}$ a.s. Thus $X(t) \in \mathcal{X}$ for a.e. t , a.s. Now, each $\mathcal{J}\hat{Z}_i^n$ is nonnegative, nondecreasing and convex, hence so is IZ_i . Therefore Z_i is nonnegative and nondecreasing. As for $\theta \cdot Y$, note that it is a.e. equal to the pathwise left-derivative of the process $\theta \cdot IY$, which, for reasons as above, has convex sample paths a.s. Hence $\theta \cdot Y$ is a.e. equal to a nondecreasing process. This shows that $(X, Y, Z) \in \mathcal{A}^P(x_0)$ and completes the proof. \square

4 A nearly optimal policy in the case of a rectangle

In this section we consider the case of a rectangular domain, where each customer class has a dedicated buffer. We have introduced in Section 2.3 some notation for this case, and identified the curve γ . In particular, the domain \mathcal{X} is given by (45), where $b_i > 0$ are fixed constants, and the parameter \mathbf{x} is given by $\theta \cdot b$. The classes are labeled so that

$$h_1\mu_1 \geq h_2\mu_2 \geq \dots \geq h_I\mu_I,$$

and, given $w \in [0, \mathbf{x}]$, $(j, \xi) = (j, \xi)(w)$ are determined by

$$w \in [\hat{b}_j, \hat{b}_{j-1})$$

and $\xi = \xi(w) = (w - \hat{b}_j)/\theta_j$, where $\hat{b}_j = \sum_{i=j+1}^I \theta_i b_i$, $j \in \{0, \dots, I\}$ and one has $0 = \hat{b}_I < \hat{b}_{I-1} < \dots < \hat{b}_1 < \hat{b}_0 = \theta \cdot b = \mathbf{x}$. With this notation, γ is given (as in (49)) by

$$\gamma(w) = \sum_{i=j+1}^I b_i e^{(i)} + \xi e^{(j)}.$$

The difficulty in treating the queueing model according to the BCP solution, as described in terms of γ , is that this curve lies along the boundary of \mathcal{X} , in particular, along the part $\partial^+ \mathcal{X} := \{x \in \mathcal{X} : x_i = b_i \text{ for some } i\}$ of the boundary $\partial \mathcal{X}$. This part corresponds to states at which some of the buffers are full. This sets up contradictory goals of keeping some of the buffers (nearly) full and at the same time avoiding any rejections except when the workload process reaches the level \mathbf{x}^* . The policy we propose is based on an approximation of γ by another curve that is bounded away from the buffer limit boundary.

Let $\varepsilon \in (0, \min_i b_i)$ be given. Let $a_i = b_i - \varepsilon$, $i \in \mathcal{I}$, and $a^* := \mathbf{x}^* \wedge (\theta \cdot a) < \mathbf{x} = \theta \cdot b$. Note that if ε is small then $a^* = \mathbf{x}^*$ (unless $\mathbf{x}^* = \mathbf{x}$). We define an approximation $\gamma^a : [0, \mathbf{x}] \rightarrow \mathcal{X}$ of γ by first defining it on $[0, \theta \cdot a]$ as the function obtained upon replacing the parameters (b_i) by (a_i) in (48) and (49). That is, for $w \in [0, \theta \cdot a]$, the variables $j = j(w)$ and $\xi = \xi(w)$ are determined via

$$w = \sum_{i=j+1}^I \theta_i a_i + \theta_j \xi, \quad j \in \{1, 2, \dots, I\}, \xi \in [0, a_j), \quad (56)$$




Figure 6: An example with imaginary buffer sizes $a_i = 1$ and reciprocal service rates $\theta_i = 1$, $i = 1, \dots, 6$. The figures depict possible states $\hat{X}^n(t) = x$ at a time when the normalized workload $\theta \cdot x = x_1 + \dots + x_6$ is around 3.5. The target population distribution is then $\gamma(3.5) = (0, 0, 0.5, 1, 1, 1)$. The class of low priority (L) is the maximal i with $x_i < a_i$. The classes served (S) are the high priority classes having positive population. Thus, for all i , i is being served provided that x_i exceeds the target population $\gamma_i(3.5)$.

and

$$\gamma^a(w) = \sum_{i=j+1}^I a_i e^{(i)} + \xi e^{(j)}. \quad (57)$$

Given $w \in [0, \theta \cdot a)$, we will sometimes refer to the unique pair (j, ξ) alluded to above as the *representation* (j, ξ) of w via (56). Next, on $[\theta \cdot a, \theta \cdot b]$ we only need the function γ^a to be continuous and satisfy the relation $\theta \cdot \gamma^a(w) = w$. For concreteness we may define it as the linear interpolation between the points $(\theta \cdot a, a)$ and $(\theta \cdot b, b)$:

$$\gamma^a(w) = a + \frac{w - \theta \cdot a}{\theta \cdot b - \theta \cdot a} (b - a), \quad w \in [\theta \cdot a, \theta \cdot b].$$

We also define $\hat{a}_j = \sum_{i=j+1}^I \theta_i a_i$, $j \in \{0, 1, \dots, I\}$, similarly to \hat{b}_j .

The definition of the policy is provided by specifying $(Z^n(t), B^n(t))$ as a function of $X^n(t)$.

Rejection policy: As under any policy, in order to keep the buffer size constraint (12), all forced rejections take place. That is, if a class- i arrival occurs at a time t when $\hat{X}_i^n(t-) + n^{-1/2} > b_i$, then it is rejected. Apart from that, no rejections occur from any class except class i^* , and no rejections occur (from any class) when $\theta \cdot \hat{X}^n < a^*$. When $\theta \cdot \hat{X}^n \geq a^*$, all class- i^* arrivals are rejected.

Service policy: For each $x \in \mathcal{X}$ define the class of low priority

$$\mathcal{L}(x) = \max\{i : x_i < a_i\},$$

provided $x_i < a_i$ for some i , and set $\mathcal{L}(x) = I$ otherwise. The complement set is the set of high priority classes:

$$\mathcal{H}(x) := \mathcal{I} \setminus \{\mathcal{L}(x)\}.$$

When there is at least one class among $\mathcal{H}(x)$ having at least one customer in the system, $\mathcal{L}(x)$ receives no service, and all classes within $\mathcal{H}(x)$, having at least one customer, receive service at a

fraction proportional to their traffic intensities. More formally, denote $\mathcal{H}^+(x) = \{i \in \mathcal{H}(x) : x_i > 0\}$, and define $\rho'(x) \in \mathbb{R}^I$ as

$$\rho'_i(x) = \begin{cases} 0, & \text{if } x = 0, \\ \frac{\rho_i 1_{\{i \in \mathcal{H}^+(x)\}}}{\sum_{k \in \mathcal{H}^+(x)} \rho_k}, & \text{if } \mathcal{H}^+(x) \neq \emptyset, \\ e^{(I)}, & \text{if } x_i = 0 \text{ for all } i < I \text{ and } x_I > 0. \end{cases} \quad (58)$$

(Note that $\mathcal{H}^+(x) = \emptyset$ can only happen if $x_i = 0$ for all $i < I$, which is covered by the first and last cases in the above display). Then for each t ,

$$B^n(t) = \rho'(\hat{X}^n(t)). \quad (59)$$

Note that when $\mathcal{H}^+(x) \neq \emptyset$,

$$\rho'_i(x) > \rho_i \quad \text{for all } i \in \mathcal{H}^+(x). \quad (60)$$

That is, all prioritized classes receive a fraction of effort strictly greater than the respective traffic intensity. Also note that $\sum_i B_i^n = 1$ whenever \hat{X}^n is nonzero. This is therefore a work conserving policy. See Figure 3 for an example of how the class with low priority and the served classes are determined.

Remark 4.1. *The only properties from the structure (58)–(59) that are actually used in the proof are (60) and $\sum_i B_i^n = 1$ if $\hat{X}^n \neq 0$; in other words we could have allowed other choices of ρ' as long as (60) and the work conserving property hold.*

Remark 4.2. *Although we have assumed $\varepsilon > 0$, the policy is well-defined even for $\varepsilon = 0$, in which case $a = b$, $a^* = \mathbf{x}^*$ and $\gamma^a = \gamma$. This policy is not used here but it is used in the next section.*

Arguing by induction on the times when the driving processes A^n and S^n jump, it is clear that there exists a unique solution to the set of equations (7)–(10), (59) along with the verbal description of the rejection mechanism. Thus the policy is well-defined.

Theorem 4.1. *For each $\varepsilon > 0$ and n , denote the policy constructed above by $U^n(\varepsilon)$. Then $\limsup_{n \rightarrow \infty} J^n(U^n(\varepsilon)) \leq V(x_0) + \alpha(\varepsilon)$, where $\alpha(\varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$.*

Remark 4.3. *i. By a usual diagonalization argument one can extract from $(U^n(\varepsilon), \varepsilon)$ a sequence U^n that is asymptotically optimal, i.e., $\limsup_{n \rightarrow \infty} J^n(U^n) = V(x_0)$.*

ii. The combination of Theorems 3.1 and 4.1 gives $\underline{V} = V(x_0)$.

Proof. We fix ε and write $U^n = (Z^n, B^n)$ for $U^n(\varepsilon)$. We denote by τ^n the time of the first forced rejection. A crucial point about the proof idea is that most of the analysis is performed on the processes up to the first forced rejection. It is established that the target state is asymptotically achieved by the proposed policy, in the sense of weak convergence as $n \rightarrow \infty$. This is done in two steps: First, the workload process $\theta^n \cdot \hat{X}^n$ is shown to converge to a RBM, and then it is shown that \hat{X}^n lies close to the minimizing curve at all times. Once these elements are established, it follows that in any finite time, τ^n is not reached, and as a result one has that (i) only rejections from class i^* occur, and only when $\theta^n \cdot \hat{X}^n \approx a^*$; (ii) the running cost is minimized locally. These elements are then combined with some integrability conditions at the last step of the proof.

We begin with the case where the system starts with initial condition close to the minimizing curve. More precisely,

$$\hat{X}^n(0) - \gamma^a(\theta \cdot \hat{X}^n(0)) \rightarrow 0 \text{ as } n \rightarrow \infty, \text{ and } \theta^n \cdot \hat{X}^n(0) \in [0, a^*] \text{ for all } n \text{ large.} \quad (61)$$

At the last step of the proof we relax this assumption.

Step 1. C-tightness for the workload and related processes. We multiply equation (18) by the vector $\theta^n = (1/\mu_i^n)_{i \in \mathcal{I}}$ and denote

$$W^{\#,n} = \theta^n \cdot \hat{W}^n, \quad X^{\#,n} = \theta^n \cdot \hat{X}^n, \quad Y^{\#,n} = \theta^n \cdot \hat{Y}^n, \quad Z^{\#,n} = \theta^n \cdot \hat{Z}^n. \quad (62)$$

We have

$$X^{\#,n} = X^{\#,n}(0) + W^{\#,n} + Y^{\#,n} - Z^{\#,n}. \quad (63)$$

Let $W^{\circ,n} := W^{\#,n}(\cdot \wedge \tau^n)$ denote the process $W^{\#,n}$ when stopped at the time τ^n . Define similarly $X^{\circ,n}$, $Y^{\circ,n}$ and $Z^{\circ,n}$. Our goal in the step is to show that the sequence $(W^{\circ,n}, X^{\circ,n}, Y^{\circ,n}, Z^{\circ,n})$ is C -tight, and that any subsequential limit $(\tilde{W}, \tilde{X}, \tilde{Y}, \tilde{Z})$ satisfies a.s.,

$$(\tilde{X}, \tilde{Y}, \tilde{Z}) = \Gamma_{[0, a^*]}[\bar{x}_0 + \tilde{W}]. \quad (64)$$

To this end, note first that the argument for C -tightness of the processes \hat{W}^n , given in the proof of the lower bound, is valid here. As a result, $W^{\#,n}$ are C -tight. Hence so are $W^{\circ,n}$.

By construction (see (59)), the policy is work conserving, namely $\sum_i B_i^n(t) = 1$ whenever $\hat{X}^n(t)$ is nonzero. By the relations (8) and (21), it follows that the nondecreasing process $\hat{Y}^{\#,n}$ does not increase when $X^{\#,n} > 0$. A similar property then holds for the stopped processes, and this can be expressed as

$$\int 1_{\{X^{\circ,n}(t) > 0\}} dY^{\circ,n}(t) = 0. \quad (65)$$

Fix $T > 0$. We show next that, as $n \rightarrow \infty$,

$$\left(\sup_{t \in [0, T]} X^{\circ,n}(t) - a^* \right)^+ \Rightarrow 0. \quad (66)$$

For $\varepsilon' > 0$ consider the event $\Omega_1^n := \{\sup_{t \in [0, T]} X^{\circ,n}(t) > a^* + \varepsilon'\}$. On this event there exist random times $0 \leq \tau_1^n < \tau_2^n \leq \tau^n$ such that $X^{\#,n}(\tau_1^n) \leq a^* + \varepsilon'/2$, $X^{\#,n}(\tau_2^n) \geq a^* + \varepsilon'$ and $X^{\#,n}(t) > a^*$ for all $t \in [\tau_1^n, \tau_2^n]$. Thus by (63) and the fact that $Y^{\#,n}$ does not increase on an interval where the system is not empty, denoting here and in the sequel $A[s, t] = A(t) - A(s)$ for any process A ,

$$\begin{aligned} (a^* + \varepsilon') - (a^* + \varepsilon'/2) &\leq X^{\#,n}[\tau_1^n, \tau_2^n] \\ &= W^{\#,n}[\tau_1^n, \tau_2^n] - Z^{\#,n}[\tau_1^n, \tau_2^n] \\ &= W^{\#,n}[\tau_1^n, \tau_2^n] - \frac{A_{i^*}^n[\tau_1^n, \tau_2^n]}{\sqrt{n}} \end{aligned}$$

where we used the fact that the policy rejects all class- i^* jobs when $X^{\#,n} > a^*$. Fix a sequence $r_n > 0$, $r_n \rightarrow 0$, such that $\sqrt{n}r_n \rightarrow \infty$. In case $\tau_2^n - \tau_1^n < r_n$, the above implies

$$\varepsilon'/2 \leq W^{\#,n}[\tau_1^n, \tau_2^n] \leq \bar{w}_T(W^{\#,n}; r_n).$$

In case $\tau_2^n - \tau_1^n \geq r_n$,

$$2\|W^{\#,n}\|_T \geq \frac{A_{i^*}^n[\tau_1^n, \tau_2^n]}{\sqrt{n}} = \hat{A}_{i^*}^n[\tau_1^n, \tau_2^n] + \frac{\lambda_{i^*}^n}{\sqrt{n}}(\tau_2^n - \tau_1^n) \geq -2\|\hat{A}_{i^*}^n\|_T + c\sqrt{n}r_n,$$

for some positive constant c . Combining the two cases, the C -tightness of $W^{\#,n}$ and the tightness of \hat{A}^n shows that $\mathbb{P}(\Omega_1^n) \rightarrow 0$ as $n \rightarrow \infty$. Since $\varepsilon' > 0$ is arbitrary, (66) follows.

Since rejections occur only when $X^{\circ,n} \geq a^*$, we have

$$\int 1_{\{X^{\circ,n}(t) < a^*\}} dZ^{\circ,n}(t) = 0.$$

Moreover, we can use (63) to write

$$X^{\circ,n} \wedge a^* = X^{\#,n}(0) + W^{\circ,n} + Y^{\circ,n} - Z^{\circ,n} + E^n, \quad E^n = (X^{\circ,n} \wedge a^*) - X^{\circ,n}.$$

Combining these relations with (65) shows that the defining relations of the Skorohod problem, namely (42)–(43), are valid here, implying

$$(a^* \wedge X^{\circ,n}, Y^{\circ,n}, Z^{\circ,n}) = \Gamma_{[0,a^*]}(X^{\#,n}(0) + W^{\circ,n} + E^n).$$

By (66), $E^n \Rightarrow 0$ uniformly on compacts. Recall that $W^{\circ,n}$ are C -tight. If \tilde{W} denotes a subsequential limit of it, using the continuity of $\Gamma_{[0,a^*]}$ and using (66) once again, shows that along the same subsequence, $(W^{\circ,n}, X^{\circ,n}, Y^{\circ,n}, Z^{\circ,n})$ converges, and that its limit satisfies (64), as claimed. The Skorohod map maps continuous paths starting in $[0, a^*]$ to continuous paths. Hence $(\tilde{W}, \tilde{X}, \tilde{Y}, \tilde{Z})$ have continuous paths a.s. This proves the claimed C -tightness of these processes.

Step 2. State space collapse. The next major step is to show that the multidimensional process \hat{X}^n lies close to the minimizing curve. More precisely, we will show that, as $n \rightarrow \infty$,

$$\Delta^n(t) := \hat{X}^n(t) - \gamma^a(X^{\#,n}(t)) \Rightarrow 0, \tag{67}$$

uniformly on compacts.

Denote by $\mathcal{G} = \{x \in \mathcal{X} : \theta \cdot x \leq a^*, x = \gamma^a(\theta \cdot x)\}$ the set of points lying on the minimizing curve, and recall the set $\partial^+ \mathcal{X} = \{x \in \mathcal{X} : x_i = b_i \text{ for some } i\}$ corresponding to the buffer limit boundary. These two compact sets do not intersect. As a result, there exists $\varepsilon_0 > 0$ such that for any $0 < \varepsilon' < \varepsilon_0$, $\mathcal{G}^{\varepsilon'}$ and $(\partial^+ \mathcal{X})^{\varepsilon'}$ do not intersect, where for a set $A \in \mathbb{R}^I$ we denote

$$A^{\varepsilon'} = \{x : \text{dist}(x, A) \leq \varepsilon'\}.$$

In what follows, it is always assumed that $\varepsilon' < \varepsilon_0$. Forced rejections occur only at times when \hat{X}^n lies in $(\partial^+ \mathcal{X})^{\varepsilon'}$ (for all n large). As a result, as long as the process \hat{X}^n lies in $\mathcal{G}^{\varepsilon'}$, no forced rejections occur. This observation can be used to deduce that $\sigma^n \leq \tau^n$, where

$$\begin{aligned} \sigma^n &= \hat{\zeta}^n \wedge \zeta^n, \\ \hat{\zeta}^n &= \inf\{t : X^{\#,n} \geq a^* + \varepsilon'\}, \quad \zeta^n = \inf\{t : \max_{i \leq I} |\Delta_i^n(t)| \geq \varepsilon'\}. \end{aligned}$$

Note carefully that σ^n is not precisely given as $\inf\{t : \hat{X}^n(t) \notin \mathcal{G}^{\varepsilon'}\}$, because $X^{\#,n}$ is defined using θ^n while γ^a and \mathcal{G} are defined with θ . However, since $\theta^n \rightarrow \theta$ and \hat{X}^n remains bounded, the conclusion that $\sigma^n \leq \tau^n$, provided that n is sufficiently large, is valid.

We turn to proving (67). It suffices to show that $\mathbb{P}(\sigma^n < T) \rightarrow 0$, for any small $\varepsilon' > 0$ and any T . Fix ε' and T . Thanks to the fact that $\sigma^n \leq \tau^n$,

$$\begin{aligned} \mathbb{P}(\sigma^n < T) &= \mathbb{P}(\sigma^n < T, \sigma^n \leq \tau^n) \\ &\leq \mathbb{P}(\hat{\zeta}^n \wedge \zeta^n \leq T \wedge \tau^n) \\ &\leq \mathbb{P}(\hat{\zeta}^n \leq T \wedge \tau^n) + \mathbb{P}(\zeta^n \leq T \wedge \tau^n). \end{aligned} \quad (68)$$

We have established in Step 1 the convergence (66), from which it follows that $\mathbb{P}(\hat{\zeta}^n \leq T \wedge \tau^n) \rightarrow 0$ as $n \rightarrow \infty$. It therefore suffices to prove the following.

Lemma 4.1. $\mathbb{P}(\zeta^n \leq T \wedge \tau^n) \rightarrow 0$ as $n \rightarrow \infty$.

Proof. On $\zeta^n \leq T \wedge \tau^n$ let $x^n := X^{\#,n}(\zeta^n) = X^{\circ,n}(\zeta^n)$ and let $j = j^n$ and ξ^n be the corresponding components from the representation (j, ξ) of x^n (with $w = x^n$).

Fix a positive integer $K = K(\varepsilon') = \lceil c_0/\varepsilon' \rceil$, where c_0 is a constant depending only on θ , whose value will be specified at a later stage of the proof. Consider the covering of $[0, \mathbf{x}]$ by the $K - 1$ intervals $\Xi_k = \mathbf{B}(k\varepsilon_1, \varepsilon_1)$, $k = 1, 2, \dots, K - 1$, where $\mathbf{B}(x, a)$ denotes $[x - a, x + a]$ and $\varepsilon_1 = \mathbf{x}/K$. Let also $\tilde{\Xi}_k = \mathbf{B}(k\varepsilon_1, 2\varepsilon_1)$.

Recall that $X^{\circ,n}$ are C -tight. Invoking the characterization of C -tightness (53)–(54), given $\delta > 0$ there exists $\delta' = \delta'(\delta, T, \varepsilon_1) > 0$ such that for all sufficiently large n ,

$$|X^{\circ,n}(s) - X^{\circ,n}(t)| \leq \varepsilon_1 \text{ for all } s, t \in [0, T], |s - t| \leq \delta', \text{ with probability at least } 1 - \delta. \quad (69)$$

Fix such δ and δ' . Then for all large n ,

$$\mathbb{P}(\zeta^n \leq T \wedge \tau^n) \leq \delta + \sum_k \mathbb{P}(\Omega^{n,k}), \quad (70)$$

where, denoting by \mathbf{T}^n the interval $[(\zeta^n - \delta' \vee 0), \zeta^n]$,

$$\Omega^{n,k} = \{\zeta^n \leq T \wedge \tau^n, x^n \in \Xi_k, X^{\#,n}(t) \in \tilde{\Xi}_k \text{ for all } t \in \mathbf{T}^n\}.$$

Note that we have used the identity $X^{\#,n} = X^{\circ,n}$ on $[0, \tau^n]$. We fix k and analyze $\Omega^{n,k}$, by an argument similar to (but somewhat more complicated than) that used in Step 1 to treat Ω_1^n .

The value assigned by the policy to B^n (see (59)) remains fixed as \hat{X}^n varies within any of the intervals $(\hat{a}_j, \hat{a}_{j+1})$. Aiming at showing that $\mathbb{P}(\Omega^{n,k}) \rightarrow 0$ as $n \rightarrow \infty$, for each k , we will first consider the case where \hat{X}^n remains in one of these intervals during the time window \mathbf{T}^n ; that is,

- (I) $\tilde{\Xi}_k \subset (0, a^*)$ and for all j , $\hat{a}_j \notin \tilde{\Xi}_k$. Then we consider the cases
- (II) $\tilde{\Xi}_k \subset (0, a^*)$ but $\hat{a}_j \in \tilde{\Xi}_k$ for some $j \in \{1, 2, \dots, I - 1\}$.
- (III) $0 \in \tilde{\Xi}_k$.
- (IV) $a^* \in \tilde{\Xi}_k$.

There may be additional intervals $\tilde{\Xi}_k$, but they are all subsets of (a^*, ∞) and therefore not important for our purpose.

(I) $\tilde{\Xi}_k \subset (0, a^*)$ and for all j , $\hat{a}_j \notin \tilde{\Xi}_k$. Note that this means that all points x in $\tilde{\Xi}_k$ lead to the same j in the representation (j, ξ) of x given by (56). Note that $j = j(k)$ depends on k only, and in particular does not vary with n . Also, $j = j^n$ under $\Omega^{n,k}$. In what follows, $j = j(k)$.

Fix $i \in \{j+1, \dots, I\}$ (unless $i = I$). We estimate the probability that, on $\Omega^{n,k}$, $\zeta^n \leq T \wedge \tau^n$ occurs by having $\Delta^n(\zeta^n) \geq \varepsilon'$. More precisely, note that $\gamma_i^a(x^n) = a_i$ (because $i > j$). Then we will show that

$$\text{for every } \varepsilon'' \in (0, \varepsilon'), \quad \mathbb{P}(\Omega^{n,k} \cap \{\hat{X}_i^n(\zeta^n) > a_i + \varepsilon''\}) \rightarrow 0 \quad \text{as } n \rightarrow \infty. \quad (71)$$

Note that γ^a is continuous and that $\Delta^n(0) \rightarrow 0$ as $n \rightarrow \infty$, by (61). Using the fact that the jumps of \hat{X}^n are of size $n^{-1/2}$, on the event indicated in (71) there must exist $\eta^n \in [0, \zeta^n]$ with the properties that

$$\hat{X}_i^n(\eta^n) < a_i + \varepsilon''/2, \quad X_i^n(t) > a_i \text{ for all } t \in [\eta^n, \zeta^n]. \quad (72)$$

On this event, during the time interval $[\eta^n, \zeta^n]$, i is always a member of $\mathcal{H}(\hat{X}^n)$, and therefore by (59)–(60), $B_i^n(t) = \rho'_i(\hat{X}^n(t)) > \rho_i + c$, for some constant $c > 0$. Thus by (21), $\frac{d}{dt}\hat{Y}_i^n \leq -\frac{\mu_i^n}{\sqrt{n}}c$. Moreover, if we define $\hat{\eta}^n = \eta^n \vee (\zeta^n - \delta')$ then for all $t \in [\hat{\eta}^n, \zeta^n]$ one has $\hat{X}^n(t) \in \tilde{\Xi}_k \subset (0, a^*)$ and therefore no rejections occur. Using these facts in (18), we have

$$\hat{X}_i^n[\hat{\eta}^n, \zeta^n] = \hat{W}_i^n[\hat{\eta}^n, \zeta^n] - c\frac{\mu_i^n}{\sqrt{n}}(\zeta^n - \hat{\eta}^n). \quad (73)$$

Again, fix a sequence $r_n > 0$ with $r_n \rightarrow 0$ and $r_n\sqrt{n} \rightarrow \infty$. If $\zeta^n - \eta^n < r_n$ and n is sufficiently large then $\hat{\eta}^n = \eta^n$, thus by (71) and the definition of η^n , $\hat{X}_i^n[\hat{\eta}^n, \zeta^n] \geq \varepsilon''/2$. As a result,

$$\bar{w}_T(\hat{W}_i^n; r_n) \geq \hat{W}_i^n[\eta^n, \zeta^n] \geq \varepsilon''/2$$

must hold. If, on the other hand, $\zeta^n - \eta^n \geq r_n$ then by (73),

$$2\|\hat{W}_i^n\|_T \geq \hat{W}_i^n[\hat{\eta}^n, \zeta^n] \geq c\frac{\mu_i^n}{\sqrt{n}}r_n \geq cr_n\sqrt{n},$$

for some constant $c > 0$. Hence the probability in (71) is bounded by

$$\mathbb{P}(\bar{w}_T(\hat{W}_i^n; r_n) \geq \varepsilon''/2) + \mathbb{P}(2\|\hat{W}_i^n\|_T \geq cr_n\sqrt{n}), \quad (74)$$

which converges to zero as $n \rightarrow \infty$, by C -tightness of \hat{W}^n . This proves (71).

Next, if we fix $i < j$ (provided $j \neq 1$) then whenever $\hat{X}_i^n > 0$, i is a member of the high priority set $\mathcal{H}(\hat{X}^n)$. Hence the same argument gives

$$\text{for every } \varepsilon'' \in (0, \varepsilon'), \quad \mathbb{P}(\Omega^{n,k} \cap \{\hat{X}_i^n(\zeta^n) > \varepsilon''\}) \rightarrow 0 \quad \text{as } n \rightarrow \infty. \quad (75)$$

Consider now j itself. We will show, for the case $j < I$,

$$\text{for every } \varepsilon'' \in (0, \varepsilon'), \quad \mathbb{P}(\Omega^{n,k} \cap \{\Delta_j^n(\zeta^n) > \varepsilon''\}) \rightarrow 0 \quad \text{as } n \rightarrow \infty. \quad (76)$$

Suppose that we show (except in the case $j = I$) for any fixed ε'' and all large n that on the event indicated in (76),

$$j \in \mathcal{H}(\hat{X}^n(t)) \text{ whenever, prior to } \zeta^n, \text{ one has } \Delta_j^n(t) \in (\varepsilon''/2, \varepsilon''). \quad (77)$$

Then we can argue as in the case of $i > j$, with the following modifications. Let $C^n(t) = \gamma_j^a(X^{\circ,n}(t))$. Then $\Delta_j^n = \hat{X}_j^n - C^n$, and similarly to (72), there exists $\eta^n \leq \zeta^n$ such that

$$\Delta_j^n(\eta^n) < \varepsilon''/2, \quad \Delta_j^n(t) > 0 \text{ for all } t \in [\eta^n, \zeta^n].$$

Since by (77) j is high priority during this interval we will still have identity (73) valid. Arguing separately for the cases $\zeta^n - \eta^n < r_n$ and $\zeta^n - \eta^n \geq r_n$, leads, in analogy to (74), to the conclusion that the probability in (76) is bounded by

$$\mathbb{P}(\bar{w}_T(\hat{W}_i^n; r_n) + \bar{w}_T(C^n; r_n) \geq \varepsilon''/2) + \mathbb{P}(2\|\hat{W}_i^n\|_T + 2\|C^n\|_T \geq cr_n\sqrt{n}). \quad (78)$$

In addition to the C -tightness of \hat{W} , we now invoke that of C^n , which follows from the continuity of γ^a and the C -tightness of $X^{\circ,n}$. This shows (74).

Now, since $\theta \cdot \gamma^a(\theta \cdot x) = \theta \cdot x$ for all $x \in \mathcal{X}$, $\theta^n \rightarrow \theta$, and γ^a uniformly continuous and \mathcal{X} bounded, we have

$$q_n := \sup_{x \in \mathcal{X}} |\theta \cdot \gamma^a(\theta^n \cdot x) - \theta \cdot x| \rightarrow 0, \quad \text{as } n \rightarrow \infty. \quad (79)$$

To show that (77) holds (except in the case $j = I$), note by (79) that $|\theta \cdot \hat{X}^n(t) - \theta \cdot \gamma^a(X^{\circ,n}(t))| \leq q_n \rightarrow 0$. If $\Delta_j^n(t) \geq \varepsilon''/2$ then

$$-\theta_j \varepsilon''/2 \geq \sum_{i \neq j} \theta_i (\hat{X}_i^n - \gamma_i^a) - \|\theta\| q_n \geq \sum_{i > j} \theta_i (\hat{X}_i^n - a_i) - \|\theta\| q_n,$$

where we used $\gamma_i^a = \gamma_i^a(X^{\circ,n}) = 0$ for $i < j$ and $\gamma_i^a = a_i$ for $i > j$. For all large n , this implies $\hat{X}_i^n < a_i$ for at least one $i > j$, by which $j \in \mathcal{H}(\hat{X}^n)$.

We can now show that $\mathbb{P}(\Omega^{n,k}) \rightarrow 0$ as $n \rightarrow \infty$. Indeed, in the case $j = I$, we have by (75), using $\gamma_i^a = 0$, $\mathbb{P}(\Omega^{n,k} \cap \{\max_{i < I} |\Delta_i^n(\zeta^n)| > \varepsilon''\}) \rightarrow 0$. By (79), $|\theta \cdot \Delta^n(\zeta^n)| \leq q_n$. Since $\theta \in (0, \infty)^I$ and $q_n \rightarrow 0$, this shows that

$$\mathbb{P}(\Omega^{n,k} \cap \{\max_{i \leq I} |\Delta_i^n(\zeta^n)| > \varepsilon''\}) \rightarrow 0. \quad (80)$$

In the case $j < I$, combining (71), (75), (76), we have $\mathbb{P}(\Omega^{n,k} \cap \{\max_{i \leq I} \Delta_i^n(\zeta^n) > \varepsilon''\}) \rightarrow 0$. Using again the fact $|\theta \cdot \Delta^n(\zeta^n)| \leq q_n \rightarrow 0$ gives that (80) is valid in this case as well.

Since ε'' is arbitrarily small, it follows from the definition of ζ^n that $\mathbb{P}(\Omega^{n,k}) \rightarrow 0$ as $n \rightarrow \infty$.

(II) $\tilde{\Xi}_k \subset (0, a^*)$ but $\hat{a}_j \in \tilde{\Xi}_k$ for some $j \in \{1, 2, \dots, I-1\}$. Let $(j^n(t), \xi^n(t))$ denote the representation (56) for $X^{\#,n}(t)$. The difficulty here is that in the time window \mathbf{T}^n , j^n varies between two values, namely j and $j+1$, and it is no longer true that $\gamma_{j+1}^a(X^{\#,n}) = a_{j+1}$ on that time interval. The way we treat this is by bounding Δ^n from above by a quantity that depends on ε_1 , rather than by an arbitrarily small ε'' . To this end, let us show that on $\Omega^{n,k}$,

$$\gamma_{j+1}^a(X^{\#,n}(t)) \geq a_{j+1} - c_1 \varepsilon_1, \quad t \in \mathbf{T}^n, \quad (81)$$

where $c_1 = 4/\theta_{\min}$ and $\theta_{\min} = \min_i \theta_i$. Indeed, we have for any $w \in \tilde{\Xi}_k$, $|w - \hat{a}_j| \leq 4\varepsilon_1$, since \hat{a}_j is also in $\tilde{\Xi}_k$. Now, if $w \geq \hat{a}_j$ then $\gamma_{j+1}^a(w) = a_{j+1}$. Otherwise,

$$w = \hat{a}_{j+1} + \theta_{j+1} \xi = \hat{a}_j - \theta_{j+1} a_{j+1} + \theta_{j+1} \xi,$$

thus $|a_{j+1} - \xi| \leq 4\theta_{j+1}^{-1}\varepsilon_1$, whence follows (81).

Now, (71) is valid for all $i > j + 1$, by the proof given in case (I). For $i = j + 1$ it is also valid, even though $\gamma_{j+1}^a(X^{\#,n}(t))$ is not necessarily equal to a_{j+1} . For $i < j$, (75) is valid with the same proof. As for $i = j$, (76) is valid with same proof (the fact that γ_j^a may assume the value zero does not affect this proof).

Combining all the estimates except for $i = j + 1$ gives, for all small ε'' ,

$$\mathbb{P}(\Omega^{n,k} \cap \{\max_{i \neq j+1} \Delta_i^n(\zeta^n) > \varepsilon''\}) \rightarrow 0.$$

For $i = j + 1$, the estimate (71) and the bound (81) give $\mathbb{P}(\Omega^{n,k} \cap \{\Delta_{j+1}^n(\zeta^n) > 2c_1\varepsilon_1\}) \rightarrow 0$ as $n \rightarrow \infty$. Along with $|\theta \cdot \Delta^n(\zeta^n)| \leq q_n$, this gives

$$\mathbb{P}(\Omega^{n,k} \cap \{\max_{i \leq I} |\Delta_i^n(\zeta^n)| > 3c_1\varepsilon_1\}) \rightarrow 0,$$

as $n \rightarrow \infty$. We now determine the constant c_0 used to define K . We do so in such a way that $3c_1\varepsilon_1 < \varepsilon'/2$. In particular, any constant $c_0 > 6c_1\mathbf{x} = 24\mathbf{x}/\theta_{\min}$ will do. This way we obtain $\mathbb{P}(\Omega^{n,k}) \rightarrow 0$ as $n \rightarrow \infty$.

(III) $0 \in \tilde{\Xi}_k$. The only difference of this case from case (I) is that during \mathbf{T}^n , \hat{X}^n may hit zero, and so by (58) and (59), B^n will be zero. However, the analysis in case (I) is performed only on intervals where $\hat{X}^n \neq 0$, and as a result gives rise to the same conclusion, namely $\mathbb{P}(\Omega^{n,k}) \rightarrow 0$ as $n \rightarrow \infty$.

(IV) $a^* \in \tilde{\Xi}_k$. In this case, during \mathbf{T}^n , $\theta \cdot \hat{X}^n$ may exceed a^* , and so rejections of class i^* customers may occur. The only way it affects the proof of case (I) is by adding a negative term to the r.h.s. of (73). However, the consequences of (73) remain valid with this addition. (Note that for all sufficiently small ε one has $\hat{a}_i \neq a^*$ for all i , hence assuming ε is sufficiently small, we do not need to check case (II) here.)

Having shown that $\mathbb{P}(\Omega^{n,k}) \rightarrow 0$ in all cases, using (70) and the fact that $\delta > 0$ is arbitrary completes the proof of the lemma. \square

As a consequence of the lemma and (68), we have $\mathbb{P}(\sigma^n < T) \rightarrow 0$ as $n \rightarrow \infty$. Since ε' is arbitrary, (67) is established.

Step 3. Weak convergence. Having shown that $\mathbb{P}(\sigma^n < T) \rightarrow 0$, we have, using $\sigma^n \leq \tau^n$, that $\mathbb{P}(\tau^n < T) \rightarrow 0$. As a result, the conclusion of Step 1 regarding the stopped processes holds also for the unstopped ones. That is, $(W^{\#,n}, X^{\#,n}, Y^{\#,n}, Z^{\#,n})$ are C -tight, and any subsequential limit satisfies (64) a.s.

Let W be an (m, σ) -BM (of dimension I) and set $\bar{W} = \theta \cdot W$. Denote by $(\bar{X}, \bar{Y}, \bar{Z})$ the triple from Proposition 2.2, i.e., $(\bar{X}, \bar{Y}, \bar{Z}) = \Gamma_{[0, \mathbf{x}^*]}(\bar{x}_0 + \bar{W})$ (note that \bar{W} is a $(\bar{m}, \bar{\sigma})$ -BM). Also let $(\bar{X}^a, \bar{Y}^a, \bar{Z}^a) = \Gamma_{[0, a^*]}(\bar{x}_0 + \bar{W})$.

For any finite T , the sequence $Z^{\#,n}(T)$ is tight. On the event $\tau^n > T$, which has overwhelming probability,

$$\hat{Z}^n(T) = \hat{Z}_{i^*}^n(T)e^{(i^*)}, \quad (82)$$

hence $\|\hat{Z}^n(T)\|$ is a tight sequence. The bound (55) thus gives the tightness of $\|\hat{Y}^n(T)\|$. The argument from the lower bound in the paragraph following (55) shows that $\hat{W}^n \Rightarrow W$ as $n \rightarrow \infty$. Thus

(64) determines the limit of the one-dimensional processes, namely $(W^{\#,n}, X^{\#,n}, Y^{\#,n}, Z^{\#,n}) \Rightarrow (\bar{W}, \bar{X}^a, \bar{Y}^a, \bar{Z}^a)$. Moreover, $(\hat{W}^n, \hat{X}^n, \hat{Y}^n, \hat{Z}^n) \Rightarrow (W, X, Y, Z)$ where $\theta \cdot X = \bar{X}^a$ and $\gamma^a(\bar{X}^a)$ (by (67)) $Z = \zeta^* \bar{Z}^a$ (by (82)) and $Y = X - x_0 - W + Z$, by (18). We obtain precisely the relations from Proposition 2.1, except that the reflection interval is $[0, a^*]$ rather than $[0, \mathbf{x}^*]$.

We have shown that, as $n \rightarrow \infty$,

$$\int_0^\infty e^{-\alpha t} [h \cdot \hat{X}^n(t) + \alpha r \cdot \hat{Z}^n(t)] dt \Rightarrow \int_0^\infty e^{-\alpha t} [h \cdot \gamma^a(\bar{X}^a(t)) + \alpha \bar{r} \bar{Z}^a(t)] dt.$$

Step 4. Convergence of costs. Since \hat{X}^n are uniformly bounded, we immediately obtain $\mathbb{E} \int_0^\infty e^{-\alpha t} h \cdot \hat{X}^n(t) dt \rightarrow \mathbb{E} \int_0^\infty e^{-\alpha t} h \cdot \gamma^a(\bar{X}^a(t)) dt$. As for the second term, we borrow an argument from [11]. Consider the probability space $(\mathbb{R}_+ \times \Omega, \mathcal{B}(\mathbb{R}_+) \times \mathcal{F}, m \times \mathbb{P})$, where $dm = \alpha e^{-\alpha t} dt$. Then the result of the previous step can be expressed as the convergence in law,

$$r \cdot \hat{Z}^n \rightarrow \bar{r} \bar{Z}^a,$$

w.r.t. the probability measure $m \times \mathbb{P}$. Thus to obtain $\mathbb{E} \int_0^\infty e^{-\alpha t} r \cdot \hat{Z}^n(t) dt \rightarrow \mathbb{E} \int_0^\infty e^{-\alpha t} \bar{r} \bar{Z}^a(t) dt$, it suffices to show the $m \times \mathbb{P}$ -uniform integrability (UI) of $r \cdot \hat{Z}^n$. For this, it suffices that

$$\limsup_n \mathbb{E} \int_0^\infty e^{-\alpha t} \|\hat{Z}^n(t)\|^2 dt < \infty. \quad (83)$$

It is established in equation (172) of [11] that

$$\mathbb{E}[(\|\hat{S}^n\|_t)^2] \leq c(1+t) \quad (84)$$

for a constant c independent of n and t , with the same estimate holding for \hat{A}^n . In what follows, we show that we can deduce (83) from (84).

To this end, recall that rejections occur only when either $\theta \cdot \hat{X}^n \geq a^*$ or, for some i , $\hat{X}_i^n \geq a_i - n^{-1/2}$. In particular, if we let $\bar{a} = a^* \wedge \min_i (a_i \theta_i / 2)$, then using the convergence $\theta^n \rightarrow \theta$, we have, for all large n , that no rejections take place when $X^{\#,n} = \theta^n \cdot \hat{X}^n < \bar{a}$. Consider the truncated version $X^{1,n} := \bar{a} \wedge X^{\#,n}$ of $X^{\#,n}$. Then by (63),

$$X^{1,n}(t) = W^{1,n} + Y^{\#,n} - Z^{\#,n}, \quad (85)$$

where we denote

$$W^{1,n} = W^{\#,n} + E^n, \quad E^n = X^{\#,n}(0) + X^{1,n} - X^{\#,n}.$$

By the above discussion,

$$\int_0^\infty 1_{\{X^{1,n} < \bar{a}\}} dZ^{\#,n} = 0,$$

and by the work conservation property,

$$\int_0^\infty 1_{\{X^{1,n} > 0\}} dY^{\#,n} = 0.$$

Moreover, the initial value $W^{1,n}(0)$ lies in $[0, \bar{a}]$. These facts dictate that $(X^{1,n}, Y^{\#,n}, Z^{\#,n})$ solves the SP on $[0, \bar{a}]$ for $W^{1,n}$. That is, $(X^{1,n}, Y^{\#,n}, Z^{\#,n}) = \Gamma_{[0, \bar{a}]}(W^{1,n})$. It is well-known (see e.g., [32])

that $\Gamma_{[0,\bar{a}]}$ is uniformly Lipschitz in the following strong sense: There exists a constant L depending only on \bar{a} , such that for every $w_1, w_2 \in D([0, t], \mathbb{R})$, one has $\|x_1 - x_2\|_t + \|y_1 - y_2\|_t + \|z_1 - z_2\|_t \leq L\|w_1 - w_2\|_t$, where $(x_i, y_i, z_i) = \Gamma_{[0,\bar{a}]}(w_i)$, $i = 1, 2$. Since the response to 0 is $(0, 0, 0)$, it follows that

$$Z^{\#,n}(t) \leq L\|W^{1,n}\|_t \leq c(\|W^{\#,n}\|_t + 2\mathbf{x}), \quad t \geq 0,$$

where we used the bound $|E^n| \leq 2\mathbf{x}$. Since θ^n converge to $\theta \in (0, \infty)^I$, this and the definitions (62) imply

$$\|\hat{Z}^n(t)\| \leq c(1 + \|\hat{W}^n\|_t).$$

Going back to (20) and using the fact that $T_i^n(t) \leq t$ for each i , and the convergence of m_i^n (see (19)),

$$\|\hat{Z}^n(t)\| \leq c(1 + t + \|\hat{A}^n\|_t + \|\hat{S}^n\|_t), \quad t \geq 0,$$

where c is independent of n and t . Combining this with (84) gives (83). Hence follows the required UI.

We have thus proved that, with $U^n = U^n(\varepsilon)$,

$$\lim_n J^n(U^n) = \lim_n \mathbb{E} \left[\int_0^\infty e^{-\alpha t} [h \cdot \hat{X}^n(t) + \alpha r \cdot \hat{Z}^n(t)] dt \right] = \mathbb{E} \left[\int_0^\infty e^{-\alpha t} [h \cdot \gamma^a(\bar{X}^a(t)) + \bar{r} \bar{Z}^a(t)] dt \right].$$

Denoting the right member above by $V(x_0; \varepsilon)$, the result will follow once we prove that $V(x_0, \varepsilon) \rightarrow V(x_0)$ as $\varepsilon \rightarrow 0$. Now, as $\varepsilon \rightarrow 0$, one has $a \rightarrow b$, $a^* \rightarrow \mathbf{x}^*$ and $\gamma^a \rightarrow \gamma$ uniformly. Moreover, the process $(\bar{X}^a, \bar{Y}^a, \bar{Z}^a)$ converges to $(\bar{X}, \bar{Y}, \bar{Z})$ in law, as can be deduced, for example, from the explicit representation of $\Gamma_{[0,a]}$ provided in [32]. Thus

$$h \cdot \gamma^a(\bar{X}^a) + \alpha \bar{r} \bar{Z}^a \rightarrow h \cdot \gamma(\bar{X}) + \alpha \bar{r} \bar{Z} = \bar{h}(\bar{X}) + \alpha \bar{r} \bar{Z},$$

in law w.r.t. $m \times \mathbb{P}$. Now, γ^a is bounded; hence to prove the convergence

$$V(x_0, \varepsilon) \rightarrow \check{V}(x_0) := \mathbb{E} \int_0^\infty e^{-\alpha t} (\bar{h}(\bar{X}) + \alpha \bar{r} \bar{Z}) dt,$$

it suffices to show the corresponding UI, and in particular, that

$$\limsup_{a^* \rightarrow \mathbf{x}^*} \mathbb{E} \int_0^\infty e^{-\alpha t} \|\bar{Z}^a(t)\|^2 dt < \infty. \quad (86)$$

To see that (86) holds, apply Ito's formula to $(\bar{X}^a(t))^2$, use the facts $\int_0^t \bar{X}^a(s) d\bar{Y}^a(s) = 0$ and $\int_0^t \bar{X}^a(s) d\bar{Z}^a(s) = aZ^a(t)$, to get

$$Z^a(t) = \frac{1}{2a} \left\{ (\bar{X}^a(0))^2 - (\bar{X}^a(t))^2 + 2 \int_0^t \bar{X}^a(s) d\bar{W}(s) + \bar{\sigma}^2 t \right\}.$$

Since \bar{X}^a is bounded by a , (86) follows easily. As a result we have $V(x_0, \varepsilon) \rightarrow \check{V}(x_0)$ as $\varepsilon \rightarrow 0$. Using integration by parts, $\check{V}(x_0) = \mathbb{E} \int_0^\infty e^{-\alpha t} (\bar{h}(\bar{X}) dt + \bar{r} d\bar{Z}(t))$. According to Proposition 2.2, this is precisely $\bar{V}(\bar{x}_0)$, because (\bar{Y}, \bar{Z}) is optimal for $\bar{V}(\bar{x}_0)$. By Proposition 2.1, $\bar{V}(\bar{x}_0) = V(x_0)$. This proves the statement of the theorem.

Step 5. General initial condition. Finally, we relax the assumption (61) on the initial condition. Here we do not give the proof in full detail, but only a brief sketch. Let ε be given and let a, a^* be as before. Let τ_0^n denote the first time when a condition analogous to (61) holds; more precisely, let $\alpha_n > 0$, $\alpha_n \rightarrow 0$ and

$$\tau_0^n = \inf\{t : \|\hat{X}^n(t) - \gamma^a(\theta \cdot \hat{X}^n(t))\| \leq \alpha_n \text{ and } \theta^n \cdot \hat{X}^n(t) \in [0, a^*]\}.$$

The idea is to show that (i) with a suitable choice of α_n , one has $\tau_0^n \rightarrow 0$ in probability, and (ii) starting from $(\tau_0^n, \hat{X}^n(\tau_0^n))$ in place of $(0, \hat{X}^n(0))$, the arguments in all the proof can be repeated without additional effort. While (ii) follows in a straightforward manner, but notationally heavy, (i) is a consequence of similar to the proof of (67). We omit the details. \square

5 On pathwise Little's law and throughput time constraints

This section is motivated by the the work of Plambeck et al. [34], where cumulative rejection costs are minimized subject to throughput time constraints in heavy traffic. The pathwise Little's law (see [33] for its original version) and the related Reiman's snapshot principle [36] imply that throughput time, regarded as a process, multiplied by the arrival rate is asymptotically equal to the queue length process in the heavy traffic limit (Proposition 5.1 below provides a precise statement for the present model). It is thus natural to expect that the solution to the problem with finite buffers that we have given can be transformed into one where throughput times are constrained, and rejection and holding costs are incurred. Our results in this direction are only partial. The main purpose of this section is to propose this problem setting and comment on relations to the main body of the paper, leaving the main question open.

Some additional notation is necessary in order to formulate the throughput time constraint problem. Recall that A_i^n and Z_i^n denote the arrival and rejection counting processes, and let $AD_i^n = A_i^n - Z_i^n$ denote the admission counting processes. Given $t \geq 0$ and $i \in \mathcal{I}$ we denote by $AT_i^n(t)$ (where AT is mnemonic for arrival time) the first time after t when a customer of class i arrives and is admitted into the system, namely

$$AT_i^n(t) = \inf\{s > t : AD_i^n(s) > AD_i^n(t)\}, \quad t \geq 0.$$

Let $DT_i^n(t)$ denote the departure time of that customer. This process can be recovered from the other processes we have defined, as follows

$$DT_i^n(t) = \inf\{s : D_i^n(s) \geq D_i^n(AT_i^n(t)) + X_i^n(AT_i^n(t))\}. \quad (87)$$

To see this, note that the time of departure of a given customer equals the time when all customers of its class present in the system at the time of its arrival (including the given customer) have departed the system. This gives the identity

$$D_i^n(DT_i^n(t)) - D_i^n(AT_i^n(t)) = D_i^n[AT_i^n(t), DT_i^n(t)] = X_i^n(AT_i^n(t)), \quad (88)$$

by which (87) follows. (Note, by right continuity, that indeed $X_i^n(AT_i^n(t))$ equals the number of class- i customers in the system at the time of arrival $AT_i^n(t)$, which includes that customer).

The time the customer spends in the system, that we will call the *throughput time*, is given by

$$\Theta_i^n(t) = DT_i^n(t) - AT_i^n(t). \quad (89)$$

Denote a diffusion scale version of the throughput time by

$$\hat{\Theta}_i^n(t) = \sqrt{n}\Theta_i^n(t).$$

Recall the distinction between an admissible control and an admissible control satisfying the buffer constraints (Definition 2.1) and the corresponding classes $\tilde{\mathcal{U}}^n$ and \mathcal{U}^n . In this section we replace the buffer constraint by an asymptotic requirement on the throughput times. Following [34], we fix constants $d_i > 0$, $i \in \mathcal{I}$, and introduce

Definition 5.1. (Asymptotic compliance) *A sequence $\{U^n\}_{n \in \mathbb{N}}$, $U^n \in \tilde{\mathcal{U}}^n$, of admissible controls, satisfying*

$$\text{for every } T, \max_i \sup_{t \in [0, T]} (\hat{\Theta}_i^n(t) - d_i)^+ \Rightarrow 0, \quad (90)$$

*is said to be **asymptotically compliant**. Denote by \mathcal{AC} the set of all asymptotically compliant sequences of admissible controls.*

Plambeck et al. [34] study AO in presence of rejection costs. In what seems to be a natural extension of their problem to include holding costs, we consider minimizing J^n among all policies satisfying (90) instead of the buffer constraints. Thus we set

$$V^{\mathcal{AC}} = \inf_{\{U^n\} \in \mathcal{AC}} \liminf_{n \rightarrow \infty} J^n(U^n).$$

We address this problem by comparing it to the problem with buffer constraints via a conditional pathwise Little's law, presented next. Its proof is deferred to the end of the section.

Proposition 5.1. *Fix $T > 0$. Given any sequence of admissible controls $U^n \in \tilde{\mathcal{U}}^n$, if \hat{Y}^n are C -tight and $\|\hat{X}^n\|_T$ are tight then $E^n \Rightarrow 0$ uniformly over $[0, T]$, where*

$$E_i^n(t) = \hat{X}_i^n(AT_i^n(t)) - \lambda_i \hat{\Theta}_i^n(t).$$

If, in addition, \hat{X}^n are C -tight then also $\tilde{E}^n \Rightarrow 0$, where

$$\tilde{E}_i^n(t) = \hat{X}_i^n(t) - \lambda_i \hat{\Theta}_i^n(t). \quad (91)$$

This suggests that a constraint d_i on $\hat{\Theta}_i^n$ should be similar a constraint $b_i = \lambda_i d_i$ on \hat{X}^n , as in the formulation with finite buffers. Recall that V denotes the value function for the RBCP. This definition depends upon the choice of the set \mathcal{X} , which we take to be rectangular as in (45), with $b_i = \lambda_i d_i$. Recall the policies defined before, and in particular, Remark 4.2 by which the policy with $\varepsilon = 0$ makes perfect sense. Here we do not have strict buffer constraints, only the requirement to meet the asymptotic constraint (90), therefore we can work with simply $\varepsilon = 0$. We denote the resulting policy by U_*^n .

Theorem 5.1. *The sequence $\{U_*^n\}$ is asymptotically compliant. Moreover, $\limsup_{n \rightarrow \infty} J^n(U_*^n) \leq V(x_0)$.*

Proof. Step 4 in the proof of Theorem 4.1 gives the upper bound on the cost. Thus it suffices to show asymptotic compliance. Now, it follows from the proof of Theorem 4.1 that the controls under consideration satisfy the assumptions of Proposition 5.1. Using this proposition along with the fact that $\mathbb{P}(\sigma^n < T) \rightarrow 0$ (see the proof of Theorem 4.1), gives the result. \square

Conjecture 5.1. *One has $V^{\mathcal{AC}} \geq V(x_0)$.*

If the above is true then, by Theorem 5.1, $\{U_*^n\}$ are AO for the problem under consideration. One might approach the conjecture by using Proposition 5.1 to connect to the lower bound of Theorem 3.1. The difficulty here is that one must consider an arbitrary sequence of controls, and there is no guarantee that the assumptions of Proposition 5.1, particularly the C -tightness of \hat{Y}^n , hold in such generality. We are able to show a partial result.

We address only policies that give rise to state space collapse. More precisely, consider a sequence $\{U^n\}_{n \in \mathbb{N}} \in \mathcal{AC}$, and write $\{U^n\} \in \widehat{\mathcal{AC}}$ if it satisfies the following. (i) Each U^n is work conserving; (ii) for some $\hat{\mathbf{x}} \in [0, \mathbf{\hat{x}})$, rejections occur only when the scaled workload exceeds $\hat{\mathbf{x}}$, and only from one particular class (save forced rejections); and (iii) for some continuous $\hat{\gamma} : [0, \mathbf{x}] \rightarrow \mathcal{X}$ satisfying

$$\{x \in \mathcal{X} : \theta \cdot x \leq \hat{\mathbf{x}}, x = \hat{\gamma}(\theta \cdot x)\} \cap \partial^+ \mathcal{X} = \emptyset, \quad (92)$$

one has $\hat{X}^n - \hat{\gamma}(\theta \cdot \hat{X}^n) \Rightarrow 0$ as $n \rightarrow \infty$. Set

$$V^{\widehat{\mathcal{AC}}} = \inf_{\{U^n\} \in \widehat{\mathcal{AC}}} \liminf_{n \rightarrow \infty} J^n(U^n).$$

Proposition 5.2. *One has $V^{\widehat{\mathcal{AC}}} \geq V(x_0)$.*

This result is far from being satisfactory. However, it shows that the two formulations are equivalent at least for this restricted class of policies. The proof is based on various elements of the proofs of the finite buffer results.

Proof. Assume, without loss of generality, that $V^{\widehat{\mathcal{AC}}} < \infty$, and consider $\{U^n\} \in \widehat{\mathcal{AC}}$ with $\liminf J^n(U^n) < \infty$. Finiteness of this quantity gives, along the lines of the proof of Theorem 3.1, that \hat{W}^n are C -tight. The assumptions on $\{U^n\}$ as a sequence in $\widehat{\mathcal{AC}}$ imply, by arguments as in the proof of Theorem 4.1, that the one-dimensional processes $(W^{\circ,n}, X^{\circ,n}, Y^{\circ,n}, Z^{\circ,n})$ are C -tight, and any subsequential limit $(\tilde{W}, \tilde{X}, \tilde{Y}, \tilde{Z})$ satisfies a.s., (64). The state space collapse assumption, along with (92) imply that, for any T , $\mathbb{P}(\tau^n < T) \rightarrow 0$, and that the un-stopped processes $(X^{\#,n}, Y^{\#,n}, Z^{\#,n})$ as well as \hat{X}^n are C -tight. Finally, C -tightness of \hat{Z}^n follows from that of $Z^{\#,n}$ by arguments as in the same proof, and that of the processes \hat{Y}^n follows equation (18) now that we have C -tightness of all the other processes involved. This verifies the assumptions of Proposition 5.1.

As a result $\tilde{E}^n \Rightarrow 0$ (where \tilde{E}^n are defined in (91)). The asymptotic compliance of the sequence of controls along with the convergence $\tilde{E}^n \Rightarrow 0$ imply the validity of the relaxed assumption (51) under which the lower bound, Theorem 3.1, is proved. Thus we conclude from Theorem 3.1 that $\liminf J^n(U^n) \geq V(x_0)$ for any $\{U^n\} \in \widehat{\mathcal{AC}}$. \square

Proof of Proposition 5.1. This proof is close to that of Lemma A.4 in the e-companion [5] of [4]. By (88) and (89), $\hat{X}_i^n(AT_i^n(t)) = n^{-1/2}D_i^n[AT_i^n(t), AT_i^n(t) + \Theta_i^n(t)]$. Now, by (9) and (11)

$$D_i^n(t) = S_i^n(T_i^n(t)), \quad \tilde{V}_i^n(t) := \hat{S}_i^n(T_i^n(t)) = \frac{S_i^n(T_i^n(t)) - \mu_i^n T_i^n(t)}{\sqrt{n}},$$

and $n^{-1/2}D_i^n = \sqrt{n}\bar{\mu}_i^n T_i^n + \tilde{V}_i^n$. By (21), recalling that $\bar{\mu}_i^n = \mu_i^n/n$, we have

$$\sqrt{n}\bar{\mu}_i^n T_i^n(t) = \sqrt{n}\bar{\mu}_i^n \rho_i t - \hat{Y}_i^n(t).$$

Hence

$$\frac{D_i^n(t)}{\sqrt{n}} = \sqrt{n}\bar{\mu}_i^n \rho_i t - \hat{Y}_i^n(t) + \tilde{V}_i^n(t),$$

and therefore

$$\hat{X}_i^n(AT_i^n(t)) = \bar{\mu}_i^n \rho_i \hat{\Theta}_i^n(t) + e_i^{1,n}(t) - e_i^{2,n}(t), \quad (93)$$

where

$$e_i^{1,n}(t) = \tilde{V}_i^n[AT_i^n(t), AT_i^n(t) + \Theta_i^n(t)], \quad e_i^{2,n}(t) = \hat{Y}_i^n[AT_i^n(t), AT_i^n(t) + \Theta_i^n(t)].$$

Now, \tilde{V}^n is C -tight by C -tightness of \hat{S}_i^n and the uniform Lipschitz property of T_i^n . The processes \hat{Y}^n are assumed to be C -tight. Thus by the assumption on $\|\hat{X}^n\|_T$, it follows that $\|\hat{\Theta}^n\|_T = \sqrt{n}\|\Theta^n\|_T$ are tight r.v.s, and thus that $e_i^{1,n}$ and $e_i^{2,n}$ all converge to zero uniformly. Since $\bar{\mu}_i^n \rho_i \rightarrow \lambda_i$, this shows the first statement of the result. The second statement now follows by the uniform convergence of $AT_i^n(t) \rightarrow t$. \square

Acknowledgement. We would like to thank Haya Kaspı for referring us to the argument in [20] and the referees for helpful comments.

References

- [1] B. Ata. Dynamic control of a multiclass queue with thin arrival streams. *Operations Research*, 54(5):876–892, 2006.
- [2] B. Ata and S. Kumar. Heavy traffic analysis of open processing networks with complete resource pooling: asymptotic optimality of discrete review policies. *Ann. Appl. Probab.*, 15(1A):331–391, 2005.
- [3] B. Ata and T. L. Olsen. Near-optimal dynamic lead-time quotation and scheduling under convex-concave customer delay costs. *Operations Research*, 57(3):753–768, 2009.
- [4] R. Atar. A diffusion regime with nondegenerate slowdown. *Oper. Res.*, 60(2):490–500, 2012.
- [5] R. Atar. A diffusion regime with nondegenerate slowdown: Appendix. *Electronic companion*, downloadable at <http://webee.technion.ac.il/people/atar/online-appendix.pdf>, 2012.
- [6] R. Atar and A. Budhiraja. Singular control with state constraints on unbounded domain. *Ann. Probab.*, 34(5):1864–1909, 2006.
- [7] R. Atar, A. Budhiraja, and R. J. Williams. HJB equations for certain singularly controlled diffusions. *Ann. Appl. Probab.*, 17(5-6):1745–1776, 2007.
- [8] R. Atar and I. Gurvich. Scheduling parallel servers in the non-degenerate slowdown diffusion regime: Asymptotic optimality results. *Ann. Appl. Probab.*, to appear, 2013.
- [9] R. Atar, A. Mandelbaum, and M. I. Reiman. Scheduling a multi class queue with many exponential servers: asymptotic optimality in heavy traffic. *Ann. Appl. Probab.*, 14(3):1084–1134, 2004.

- [10] R. Atar and N. Solomon. Asymptotically optimal interruptible service policies for scheduling jobs in a diffusion regime with nondegenerate slowdown. *Queueing Systems Theory Appl.*, 69(217–235), 2011.
- [11] S. L. Bell and R. J. Williams. Dynamic scheduling of a system with two parallel servers in heavy traffic with resource pooling: asymptotic optimality of a threshold policy. *Ann. Appl. Probab.*, 11(3):608–649, 2001.
- [12] M. N. Bennani and D. A. Menasce. Resource allocation for autonomic data centers using analytic performance models. In *Autonomic Computing, 2005. ICAC 2005. Proceedings. Second International Conference on*, pages 229–240. IEEE, 2005.
- [13] P. Billingsley. *Convergence of probability measures*. Wiley Series in Probability and Statistics: Probability and Statistics. John Wiley & Sons Inc., New York, second edition, 1999. ISBN 0-471-19745-9. x+277 pp. A Wiley-Interscience Publication.
- [14] M. Bramson. State space collapse with application to heavy traffic limits for multiclass queueing networks. *Queueing Systems Theory Appl.*, 30(1-2):89–148, 1998.
- [15] A. Budhiraja and A. P. Ghosh. Diffusion approximations for controlled stochastic networks: an asymptotic bound for the value function. *Ann. Appl. Probab.*, 16(4):1962–2006, 2006.
- [16] A. Budhiraja and A. P. Ghosh. Controlled stochastic networks in heavy traffic: convergence of value functions. *Ann. Appl. Probab.*, 22(2):734–791, 2012.
- [17] C. Buyukkoc, P. Varaiya, and J. Walrand. The $c\mu$ rule revisited. *Adv. in Appl. Probab.*, 17(1):237–238, 1985.
- [18] D. R. Cox and W. L. Smith. *Queues*. Methuen’s Monographs on Statistical Subjects. Methuen & Co. Ltd., London; John Wiley & Sons Inc., New York, 1961. xii+180 pp.
- [19] J. G. Dai and W. Dai. A heavy traffic limit theorem for a class of open queueing networks with finite buffers. *Queueing Systems Theory Appl.*, 32(1-3):5–40, 1999.
- [20] C. Dellacherie and P.-A. Meyer. *Probabilities and potential*, volume 29 of *North-Holland Mathematics Studies*. North-Holland Publishing Co., Amsterdam-New York; North-Holland Publishing Co., Amsterdam-New York, 1978. ISBN 0-7204-0701-X. viii+189 pp.
- [21] R. Durrett. *Probability: theory and examples*. Duxbury Press, Belmont, CA, second edition, 1996. ISBN 0-534-24318-5. xiii+503 pp.
- [22] S. N. Ethier and T. G. Kurtz. *Markov processes*. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. John Wiley & Sons Inc., New York, 1986. ISBN 0-471-08186-8. x+534 pp. Characterization and convergence.
- [23] S. Ghamami and A. R. Ward. Dynamic scheduling of a two-server parallel server system with complete resource pooling and reneging in heavy traffic: Asymptotic optimality of a two-threshold policy. *Mathematics of Operations Research*, 38(4):761–824, 2013.
- [24] A. P. Ghosh and A. P. Weerasinghe. Optimal buffer size for a stochastic processing network in heavy traffic. *Queueing Syst.*, 55(3):147–159, 2007.
- [25] J. M. Harrison. Brownian models of queueing networks with heterogeneous customer populations. In *Stochastic differential systems, stochastic control theory and applications (Minneapolis, Minn., 1986)*, volume 10 of *IMA Vol. Math. Appl.*, pages 147–186. Springer, New York, 1988.
- [26] J. M. Harrison. Brownian models of open processing networks: canonical representation of workload. *Ann. Appl. Probab.*, 10(1):75–103, 2000.
- [27] J. M. Harrison. A broader view of Brownian networks. *Ann. Appl. Probab.*, 13(3):1119–1150, 2003.
- [28] J. M. Harrison and M. I. Taksar. Instantaneous control of Brownian motion. *Math. Oper. Res.*, 8(3):439–453, 1983.
- [29] J. M. Harrison and J. A. Van Mieghem. Dynamic control of Brownian networks: state space collapse and equivalent workload formulations. *Ann. Appl. Probab.*, 7(3):747–771, 1997.
- [30] J. M. Harrison and R. J. Williams. Workload reduction of a generalized Brownian network. *Ann. Appl. Probab.*, 15(4):2255–2295, 2005.
- [31] J. Jacod and A. N. Shiryaev. *Limit theorems for stochastic processes*, volume 288 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1987. ISBN 3-540-17882-1. xviii+601 pp.

- [32] L. Kruk, J. Lehoczy, K. Ramanan, and S. Shreve. An explicit formula for the Skorokhod map on $[0, a]$. *Ann. Probab.*, 35(5):1740–1768, 2007.
- [33] R. Mazumdar, V. Badrinath, F. Guillemin, and C. Rosenberg. A note on the pathwise version of Little’s formula. *Oper. Res. Lett.*, 14(1):19–24, 1993.
- [34] E. Plambeck, S. Kumar, and J. M. Harrison. A multiclass queue in heavy traffic with throughput time constraints: asymptotically optimal dynamic controls. *Queueing Syst.*, 39(1):23–54, 2001.
- [35] E. L. Plambeck. Optimal leadtime differentiation via diffusion approximations. *Oper. Res.*, 52(2):213–228, 2004.
- [36] M. I. Reiman. The heavy traffic diffusion approximation for sojourn times in Jackson networks. In *Applied probability—computer science: the interface, Vol. II (Boca Raton, Fla., 1981)*, volume 3 of *Progr. Comput. Sci.*, pages 409–421. Birkhäuser Boston, Boston, MA, 1982.
- [37] M. Rubino and B. Ata. Dynamic control of a make-to-order, parallel-server system with cancellations. *Oper. Res.*, 57(1):94–108, 2009.
- [38] M. Shifrin, R. Atar, and I. Cidon. Optimal scheduling in the hybrid-cloud. In *IFIP/IEEE International symposium on integrated network management*. IEEE, 2013, to appear.
- [39] W. E. Smith. Various optimizers for single-stage production. *Naval Res. Logist. Quart.*, 3:59–66, 1956.
- [40] J. A. van Mieghem. Dynamic scheduling with convex delay costs: the generalized $c\mu$ rule. *Ann. Appl. Probab.*, 5(3):809–833, 1995.
- [41] A. R. Ward and S. Kumar. Asymptotically optimal admission control of a queue with impatient customers. *Math. Oper. Res.*, 33(1):167–202, 2008.
- [42] W. Whitt. Weak convergence theorems for priority queues: Preemptive-resume discipline. *J. Appl. Probability*, 8:74–94, 1971.
- [43] R. J. Williams. Diffusion approximations for open multiclass queueing networks: sufficient conditions involving state space collapse. *Queueing Systems Theory Appl.*, 30(1-2):27–88, 1998.